

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 1/31/96	3. REPORT TYPE AND DATES COVERED Final Technical Report 6/1/92-10/1/95	
4. TITLE AND SUBTITLE Processes involved in the integration of pictures and discourse			5. FUNDING NUMBERS F49620-92-J-0310 61103D 3484/S4	
6. AUTHOR(S) Arthur M. Glenberg				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Psychology University of Wisconsin 1202 West Johnson Street Madison, WI 53706			8. PERFORMING ORGANIZATION REPORT NUMBER AFOSR-TR-96-00489	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NL Bolling AFB, DC 20332-8080			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Dr John F. Tangney			19960220 025	
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) There is no doubt that pictures can enhance text comprehension, and sometimes the effect is extraordinary. But how do pictures have this effect? One hypothesis is that pictures help to structure a spatial mental model. Indeed, initial results were strongly compatible with this hypothesis and contrary to others. Subsequent research was directed at characterizing the mental model. One possibility is that the mental model has a Euclidean structure and that it is formed using the spatial component of working memory. Although a type of working memory is clearly involved in mental model formation, three types of research have demonstrated that the mental model is unlikely to be Euclidean. First, the degree of priming between elements of the mental model is unrelated to Euclidean distance (see Kramer, Langston, & Glenberg); instead, functional distance seems to matter. Second, there is little benefit produced by readers producing drawings when they read unillustrated texts (see Robertson & Glenberg). Third, a careful analysis of memory and its functioning reveals that accounts based on the concept of embodiment provide a more satisfactory explanation than accounts based on considering mental models to be constructed in an unstructured Euclidean space (see Glenberg).				
14. SUBJECT TERMS reading comprehension, pictures, mental models, embodiment			15. NUMBER OF PAGES 101	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

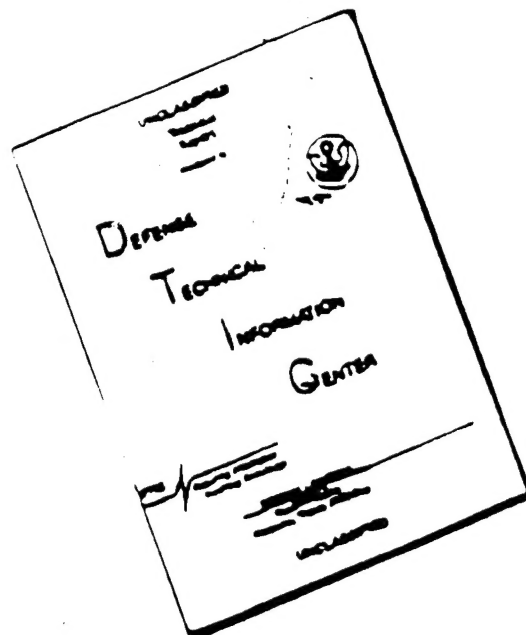
NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

REF 9

DTIC QUALITY INSPECTED 1

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

Executive Summary

The initial research conducted on this project demonstrated that pictures help text comprehension by assisting in the structuring of a mental model for the text. Later research investigated the nature of those mental models. Because that later research is not yet published, this Executive Summary and the report focus on the later research.

Langston, Kramer, and Glenberg (see page 6 for full report)

Mental models of text are representations of what the text is about (i.e., situations), rather than representations of the text itself. Many mental model theories propose that mental models, like real situations, are played out in a medium that has the properties of Euclidean space. An expectation of such a medium is that distance has functional consequences. For example, when mentally manipulating one element of the representation, other elements that are spatially close will be noticed and their short-term accessibility enhanced. To test this noticing hypothesis, subjects read texts that described the object by object construction of a spatial layout. According to the text, a critical object ended up close to a target object (in the spatial layout) or far from the target object. In neither case, however, was the relation between the critical object and the target object explicitly described in the text. The noticing hypothesis predicts that the accessibility of the target object will be enhanced when the critical object is close to it. We tested this prediction in seven experiments in which we also manipulated the number of objects described, whether the description was accompanied by a diagram, the presentation modality of the description, the number of dimensions in the spatial layout, and the measurement of accessibility. We failed to find consistent support for the noticing hypothesis. The data compel the conclusions that a) spatial representations can be formed when reading, b) these representations do not support automatic noticing of implicit spatial relations, c) it is likely that the spatial representation is more topological than Euclidean.

Robertson and Glenberg (see page 38 for full report)

Can readers enhance comprehension by drawing their own pictures for the texts they are reading? Spatial mental model theory implies the answer is yes, because these models are presumed to underlie comprehension. Participants performed one of three study tasks, summarizing, drawing, or read-only, while reading short texts. In Experiment 1, draw subjects performed better on an immediate comprehension test than summarize subjects, but worse than read-only subjects. In Experiment 2, participants were tested immediately after reading and after a 12 minute delay. The immediate test condition replicated Experiment 1, whereas the delayed test showed no difference between read-only and drawing. At a 40 minute retention interval in Experiment 4, draw subjects performed worse than read-only subjects. Experiment 3 established that these effects were not due to oddities in the stimuli. These results demonstrate a serious flaw in spatial mental model theory. We modify the theory by proposing that readers will benefit from drawing provided they have domain-specific strategies relevant to producing beneficial pictures.

Glenberg (see page 64 for full report)

Suppose that memory evolved in the service of perception and action in a three-dimensional world. As such, memory specializes in the representations that support real, physical actions involving the body and the environment. These embodied representations can be meshed the way multiple physical actions can be combined, and the idea of meshed patterns can be used to replace the theoretically empty term "association." This analogical approach to memory addresses the symbol grounding problem, automatic and effortful uses of memory, and language comprehension, in particular, mental model theory.

Cumulative list of personnel:

Arthur Glenberg (PI, faculty)
Douglas Kramer (RA, graduate student)
Peter Kruley (RA, graduate student)
William Langston (RA, graduate student)
David Robertson (RA, graduate student)

Publications:

- Glenberg, A. M., & Langston, W. E. (1992). Comprehension of illustrated text: Pictures help to build mental models. *Journal of Memory and Language*, 31, 129-151.
- Glenberg, A. M., & Kruley, P. (1992). Pictures and anaphora: Evidence for independent processes. *Memory & Cognition*, 20, 461-471.
- Glenberg, A. M., & Mathew, S. (1992). When minimalism is not enough: Mental models in reading comprehension. *PSYCOLOQUY*.92.3.64.reading-inference-2.1.glenberg-mathews.
- Glenberg, A. M., & McDaniel, A. (1992). Mental models, pictures, and text: Integration of spatial and verbal information. *Memory & Cognition*, 20, 458-460.
- Glenberg, A. M. (1993). Distributed practice effects. In L. R. Squire (Ed.), *Encyclopedia of Learning and Memory*. New York: Macmillan.
- Glenberg, A. M. (1993). Comprehension while missing the point: More on minimalism and models. *PSYCOLOQUY*.93.4.31.reading-inference.13.glenberg.
- Glenberg, A. M., Kruley, P., & Langston, W. E. (1994). Analogical processes in comprehension: Simulation of a mental model. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*.
- Strauman, T. J., & Glenberg, A. M. (1994). Self-concept and body image disturbance: Which self-beliefs predict body size overestimation? *Cognitive Therapy and Research*, 18, 105-125.
- Kruley, P., Sciamia, S. C., & Glenberg, A. M. (1994). On-line processing of textual illustrations in the visuo-spatial sketchpad: Evidence from dual-task studies. *Memory & Cognition*, 22, 261-272.
- Glenberg, A. M., & Grimes, T. (1995). Memory and faces: Pictures help you to remember who said what. *Personality and Social Psychology Bulletin*, 21, 196-206.
- Glenberg, A. M. (in press). Mental models, space, and embodied cognition. T. Ward & S. Smith (Eds.), *Creative Cognition*.

Submitted for publication:

- Glenberg, A. M. What memory is for. *Behavioral and Brain Sciences*.
- Glenberg, A. M., Schroeder, J. L., & Robertson, D. A. (submitted for publication). Remembering behavior.
- Langston, W., Kramer, D. C., & Glenberg, A. M. (submitted for publication). Spatial mental models from text are non-Euclidean.

Invited addresses:

- Analogical processes in comprehension. Colloquium at University of Chicago, November, 1992.
- Analogical processes in comprehension. Colloquium at Northwestern University, December, 1992.

- Analogical processes in comprehension. Colloquium at ZiF, University of Bielefeld, Germany, January, 1993.
- Analogical processes in comprehension. Colloquium at Department of Linguistics, University of Bielefeld, 1993.
- Analogical processes in comprehension. Colloquium at University of Giessen, Germany, February, 1993.
- Analogical processes in comprehension. Colloquium at CNRS, University of Paris-Sud, March, 1993.
- Metaphor, models, and memory. Colloquium at University Rene Descartes, Paris, March, 1993.
- Metaphor, models, and memory. Colloquium at University of Leuven, Leuven, Belgium, March 1993.
- Analogical processes in comprehension. Colloquium at Max-Planck-Institut for Psycholinguistik, Nijmegen, The Netherlands, March, 1993.
- Metaphor, models, and memory. Workshop on Mental Models, ZiF, University of Bielefeld, Germany, April, 1993.
- What are analogical components doing there? Festschrift for Walter Kintsch, University of Colorado, Boulder, CO, June 1993.
- Suppressing the environment: A possible link between prediction, memory, and language comprehension.
- Invited symposium on The role of inhibition across psychological domains. Midwestern Psychological Association, Chicago, May, 1994.
- Finding meaning in cognitive psychology. Colloquium at University of Minnesota, May, 1994.
- Mental models are not (very) spatial. European Association for Research on Learning and Instruction/Special Interest Group. Helsinki, Finland, June, 1994.
- Mental models are not (very) spatial. Lake Geneva Language Comprehension Conference, Lake Geneva, WI, July, 1994.
- Finding meaning in cognitive psychology. Colloquium at National Institutes of Health, Bethesda, MD, September, 1994.
- Mental models, space, and embodied cognition. Creative Cognition Conference, Texas A&M University, May 1995.
- Comments on Keltner, Habel, and Kaup's "Taking the functional aspect of mental models as a starting point for studying discourse comprehension." Workshop on Mental Models, ZiF, University of Bielefeld, Germany, July, 1995.
- Remembering behavior. Institute of Cognitive Science, University of Colorado, Boulder, October 1995.
- Not Propositions. Institute of Cognitive Science, University of Colorado, Boulder, December 1995.
- Finding meaning in cognitive psychology. Department of Cognitive Psychology, La Laguna University, Spain, January 1996.
- Not Propositions. Department of Cognitive Psychology, La Laguna University, Spain. January 1996.
- Suppressing the environment: A possible link between prediction, memory, and language comprehension. Department of Cognitive Psychology, La Laguna University, Spain. January 1996.

Conference Presentations:

- Glenberg, A. M. Spatial mental models. Center for Interdisciplinary Studies, University of Bielefeld, Germany, July, 1992.
- Glenberg, A. M. Evidence for functional spatial models in text comprehension. Paper presented at Tagung experimentell arbeitender Psychologen, Munster, Germany,

April, 1993.

Langston, W., Kramer, D. C., & Glenberg, A. M. Mental models are not (very) spatial. Presented at the meeting of the Psychonomic Society, St. Louis, November, 1994.

Robertson, D. A., & Glenberg, A. M. Does sketching while reading facilitate learning from texts? A comparison of three study tasks. Paper presented at the sixty-seventh annual meeting of the Midwestern Psychological Association, Chicago, May, 1995.

Robertson, D. A., & Glenberg, A. M. Sketching while reading does not facilitate learning from texts. Paper presented at the 5th Annual Meeting of the Society for Text and Discourse, Albuquerque, NM, July, 1995.

Glenberg, A. M., Schroeder, J. L., & Robertson, D. A. Remembering behavior. Presented at the meeting of the Psychonomic Society, Los Angeles, November, 1995.

Langston, Kramer, and Glenberg, full report

Spatial Mental Models from Text are Non-Euclidian

The field is reaching the consensus that an important component of language comprehension consists of constructing mental models (Johnson-Laird, 1983; Gentner & Stevens, 1983) or situational models (van Dijk & Kintsch, 1983). In essence, what is meant by a mental model is a representation of what the language (or text) is about, rather than a representation of the language itself. Other than this essence, however, there is little agreement as to how to characterize a mental model. Is it an image? A schema? A set of propositions? In this paper we take a step toward a characterization of mental models, although the outcome of that step is to rule out a possibility rather than to confirm one.

Various scientists have proposed that mental models not only capture spatial relations, but that they are constructed in a spatial medium, a medium analogous to the Euclidean, three-dimensional space of everyday experience. We begin with a brief review of the data consistent with this view, and we propose a test of it. We then present the results of seven experiments that lead us to conclude that this view is incorrect. Although the data are overwhelming in demonstrating that spatial representations are constructed, there is no support for a prediction derived from a Euclidean view. In the General Discussion we review other research that supports these conclusions and we offer a few suggestions to replace the notion that mental models are built in a Euclidean medium.

A number of theorists have claimed that spatial mental models do have Euclidean, or metric, properties. For examples, such a claim is made by Johnson-Laird (1983, page 422), and it is an important component of the computer simulation described by Glenberg, Kruley, and Langston (1994). While admitting to other possibilities, Wagener and Wender (1985) "assume that this [spatial] representation has properties that are in some sense functionally equivalent to those of an actual scene. This means that the space is continuous and distances are determined by an Euclidian metric" (page 132). Denis, Gonçlaves, and Memmi (in press) conclude that "These data support the claim that images generated from verbal descriptions can have metric properties..." In discussing accessibility of representations in a mental model encoding objects and a protagonist, Rinck and Bower (1995) imply a Euclidean representation by statements such as "The accessibility of objects should gradually decrease with increasing distance from the protagonist" (page 112).

Why would anyone believe that mental representations are analogous to space? One reason is that much of the research supporting the concept of mental models of discourse has traded on spatial relations. Garnham (1992) has discussed one reason for this: Texts that describe spatial layouts can often be constructed so that the propositional structure of the text is different from the spatial structure of the situation described by the text. That is, when objects are close in space, the words used to refer to and describe those objects need not be close (spatially, temporally, or propositionally) in the description. Thus, demonstrating an effect on comprehension of the structure of the situation indicates that it must be considered in addition to the structure of the description. According to Garnham, ease of experimental manipulation has driven at least some of the theorizing about mental models. What follows is a brief review of some articles that have taken advantage (to a greater or lesser extent) of the possibility of dissociating propositional structure of the description and the spatial structure of the situation described.

Mani and Johnson-Laird (1982) presented their subjects with descriptions of spatial layouts. The description could be determinate (correspond to only one topographical layout) or indeterminate. For the indeterminate layouts, either subjects would have to construct multiple layouts, or remember the descriptions verbatim. A major finding was that for the indeterminate layouts, subjects were more accurate verifying verbatim statements than in verifying correct but

unstated inferences. For the determinate layouts (when, supposedly a single spatial representation could be formed), subjects were equally accurate in verifying verbatim and inferential statements. Apparently, when possible to do so, people construct spatial representations (but see Payne, 1993, for an impressive failure to replicate).

Denis and Cocude (1989, 1992) have demonstrated that scanning a representation of an island produces similar results whether the scan is controlled by a mental representation formed from a linguistic description, a mental representation formed from observing a map, or a physical representation, that is, the map itself. In these experiments, subjects heard a description of a circular island with critical landmarks located by clock-face coordinates. The subjects were then asked to mentally scan from one location to another, and the time taken to do so was recorded. After hearing the text six times, the correlation between distance (on the map described by the text) and time to scan the putative mental model was close to .90, comparable to the correlation when an actual map (rather than a text) was studied but the scanning was done from memory (.90), and comparable to the correlation when a physical map was scanned (.87). These data strongly suggest that the constructed map had the spatial characteristics of the real map: that is, the constructed map is Euclidean.

Morrow, Bower, and Greenspan (1989) demonstrated another sort of spatial distance effect. In their experiments, subjects first memorized the locations of objects in a ten-room building. They then read a text describing the movements of a protagonist in the building, and the subjects were probed with pairs of objects. The task was to answer "yes" if the objects were in the same room (not necessarily the room in which the protagonist was located). Correct reaction time on this task was positively related to the distance between the protagonist and the room containing the objects. Wilson, Rinck, McNamara, Bower, and Morrow (1993) report a qualification of this distance effect, but Rinck and Bower (1995) have replicated the effect using a different dependent variable not subject to the qualification. Rinck and Bower found that the time to read a sentence referring to an object was positively correlated with the distance between the current location of the protagonist and the room containing the referred-to object. They conclude, "Taken together, the four experiments reported here provide evidence for the claim that spatial distance in the reader's situation model of a narrative influences the accessibility of referents for anaphoric expressions" (page 129). In short, the data are consistent with a Euclidean view in which spatial distance is functional.

O'Brien and Albrecht (1992) demonstrated that readers are sensitive to the location of a protagonist even without pre-memorization of locations. In these experiments subjects read sentences such as "As Kim stood inside/outside the health club she felt a little sluggish." Six lines later, the subject might encounter, "She decided to go outside and stretch her legs." Reading time of this critical sentence was significantly faster when it was consistent with the initial location of the protagonist (inside to outside), than when it was inconsistent (outside to outside).

Glenberg, Meyer, and Lindem (1987) showed effects of spatial distance between a protagonist and a critical object. In these experiments the critical object (such as "sweatshirt") was described as spatially close to the protagonist in the associated condition, as in "After doing a few warm-up exercises, [John] put on his sweatshirt and went jogging." In the dissociated condition, the critical object and the protagonist were described as spatially separated, as in "After doing a few warm-up exercises, [John] took off his sweatshirt and went jogging." Accessibility of the critical object (e.g., sweatshirt) was assessed after reading another sentence in which John was foregrounded but sweatshirt never mentioned. Accessibility of "sweatshirt" was found to be greater in the spatially close, associated condition than in the spatially separated, dissociated condition.

Finally, Glenberg and Langston (1992) proposed that under certain circumstances spatial

mental models could be used to represent non-spatial relations. In those experiments, subjects read descriptions of the temporal order of the steps in multi-step procedures, and each subject was tested on memory for the order in which the steps should be performed, not the order of the description of the steps in the text. The texts were completely explicit regarding the order in which the steps should be performed, but because two of the steps were to be performed at the same time, whereas their descriptions proceeded serially, the order of performance could be dissociated from the order of description. The data suggested a) that the structure of memory for order was affected by order of description, but b) when the texts were accompanied by diagrams in which spatial distances among the steps were analogous to temporal separations, the effect of order of description was much reduced. In the latter case, memory reflected the mental model, that is, the representation of how to perform the multi-step task, not memory for the description itself. Importantly, when a diagram was available, readers seemed to have noted temporal relations among steps spatially close in the diagram (and presumably close in the mental model based on the diagram) more readily than temporal relations among steps spatially distant in the diagram.

Glenberg and Langston used their results to support the noticing hypothesis, which sets the stage for the experiments described here. The noticing hypothesis provides one explanation for why forming mental models enhances comprehension: the spatial medium in which the model is constructed supports a type of inference making called noticing. The noticing hypothesis is based on the assumptions that a) mental models are constructed in a spatial medium, such as the limited-capacity visuo-spatial scratchpad identified by Baddeley (1986); b) text-relevant dimensions (e.g. time) can be assigned to spatial dimensions, although the assignment may require expertise in a domain or visual support such as that provided by a diagram; c) text-relevant concepts (e.g., steps in a procedure) are represented by pointers arrayed within the medium and each pointer is associated with information about the concept in memory. Furthermore, the spatial distances between pointers is cognitively meaningful. That is, pointers that are spatially close in working memory represent concepts that are more closely (or strongly) related along the text-relevant dimension assigned to the spatial dimension. d) Mental models are updated by introducing new pointers or rearranging existing pointers to reflect the situation described by the text. e) When an existing pointer is moved, or when a new pointer is introduced into the model, that pointer is attended, and any pointer within the "spotlight" of attention is noticed. "Noticing" means that the relation between the pointers is specifically encoded and added to the associated information in memory. The operation of this noticing process was formalized in a computer simulation described by Glenberg et al. (1994).

We found the noticing hypothesis appealing for several reasons. First, it provides one reason why forming a mental model facilitates comprehension: the mental model allows for a type of inference making (noticing on the basis of spatial contiguity) not easily available to a propositional system. Second, unlike many schemes for inference generation, noticing is highly constrained: It only occurs along text-relevant dimensions assigned to spatial dimensions; it is constrained by the capacity limits of working memory; it occurs only during updating a mental model and then only between spatially close pointers. Third, the simulation model bears at least a family resemblance to models developed by Johnson-Laird (e.g., Johnson-Laird, Byrne, & Tabossi, 1989) to account for reasoning using mental models. Fourth, Glenberg et al. (1994) show how a simulation model incorporating noticing can account for much of the data reviewed above. Finally, the noticing hypothesis is implied by many of the theoretical claims about the operation of spatial mental models (e.g., Denis & Cocude, 1992; Rinck & Bower, 1995). That is, hypotheses that assume a Euclidean medium for mental models assume that distance within the medium has functional consequences; items that are farther apart are less likely to interact. Thus, testing the noticing hypothesis gives us some purchase on the claim that spatial mental models are Euclidean.

The seven experiments reported here were designed to test the noticing hypothesis. The

basic logic was to have subjects read (or listen to) a text describing a spatial layout. An example text appears in Table 1. Each text described the relations between some objects explicitly, whereas the spatial relations between other objects were left implicit. A speeded recognition probe (Experiments 1-4) or a sentence reading task (Experiments 5-7) was used to assess the accessibility of the items in the layout. The primary manipulation was the distance (in the described spatial layout, not the text) between the target (i.e., probed) item and another most recently mentioned item. The most recently mentioned item is an item that is described in the text immediately before the probe is presented. In Table 1, "seaweed" is the most recently mentioned item. Sometimes the most recently mentioned item was near to the target in the spatial layout (notice condition), and sometimes it was far from the target (not-notice condition). Importantly, the spatial distance between the target and the most recently mentioned item was never described explicitly. Instead, the relation between the most recently mentioned item and some other object (the "big rock" in Table 1) controlled the relation between the target and the most recently mentioned item. In the notice condition, when the most recently mentioned item is near to the target (in the spatial layout, and hence in the putative spatial model), more noticing of the target should occur than when the two objects are separated from one another, as in the not-notice condition. We assessed the degree of noticing by measuring the accessibility of the target soon after the most recently mentioned item was introduced by the text.

Insert Table 1 about here

Table 2 summarizes the results of the seven experiments, and it provides a guide to the detailed exposition that follows. The two dependent variables that we used are listed in the DV column. One dependent variable was speeded recognition, the time to recognize the probe word (naming the target item) as an item that occurred in the text. The probe word occurred immediately after the most recently mentioned item. Speeded recognition is a commonly used measure of accessibility (e.g., Glenberg et al., 1987; de Vega, 1995). The other dependent variable was time to read an about sentence. The about sentence followed the most recently mentioned item, and the sentence referred to the target item. Hence the time to read the sentence included the time needed to make an anaphoric reference to the target item. Sentence reading time is a common measure of accessibility of the referents of anaphors (e.g., Rinck & Bower, 1995). Most of the experiments used visual presentation of the text (as indicated in the Modality column of Table 2), but we also report one experiment (Experiment 3) that used auditory presentation. The next column (Item/Dim) indicates the item number of the most recently mentioned item. Thus, across the experiments, the spatial layout of 3, 4, or 6 items was described before accessibility of the target item was assessed. The number following the slash mark indicates whether the spatial layout was in one or two dimensions¹. In most of the experiments, the distance in the spatial layout between the most recently mentioned item and the target item was confounded with number of intervening items. That is, no items intervened in the notice condition, and several items might intervene in the not-notice condition. This confound was removed in those conditions labeled "far" in Table 2. For those comparisons, we computed the noticing effect as the difference between the notice-far and the far conditions (see Experiment 2). In those conditions, the spatial distance (in the layout) between the target and the most recently mentioned item was manipulated without intervening items. The column labeled "Picture" indicates whether the verbal description of the layout was accompanied by a picture illustrating the layout of the first three items. The "unlabeled" picture condition in Experiment 2 illustrated the layout using tokens (boxes), but did not label the tokens with the item names.

Insert Table 2 about here

One index of the noticing effect is the difference between the median value of the dependent variable in the not-notice condition and the median value of the dependent variable in the notice condition. The noticing hypothesis predicts a positive difference. This index is provided in the column labeled "Notice Effect." The final four columns present statistical assessments of the generalizability of the notice effect. A noticing effect that is significantly different from zero is italicized. The first two columns report the value of the paired t-statistic comparing the medians. The "Standard" analysis eliminated individual observations if a) response to the probe was incorrect, or b) response to a comprehension question presented after the text was incorrect. The Loose analysis eliminated observations according to condition a, but not condition b. Thus, the medians computed in the Loose analysis are based on more observations. Bush, Hess, and Wolford (1993) have demonstrated that comparisons based on medians may not be statistically powerful. Hence, we also include two analyses (Standard and Loose) based on a set of transformations that Bush et al. found to be the most powerful without sacrificing protection against Type 1 errors. The transformations are intended to achieve power by a) elimination of outliers by throwing out the longest and the shortest observations in each condition for each subject, b) normalization of skewed reaction time distributions by using the logarithm of the speeded recognition or reading times, and c) elimination of between-subject variability in the size of the effect by using z-scores computed for each subject. Thus, after elimination of outliers, the individual observations were converted into logarithms. Next, a single mean and standard deviation was computed for all of a given subject's scores (regardless of condition), and the subject's mean and standard deviation were used to compute z-scores for all of the observations. Finally, a mean z-score was computed for each subject in each condition. The mean of these mean z-scores will always equal zero (so that there is no between-subject variability), but within-subject differences among conditions are still preserved. The last two columns in Table 2 present the results of the dependent sample t-test comparing the mean z-scores of the notice and not-notice conditions. For each analysis, the number of subjects entering into the analysis is noted in parentheses.

Thus, over the course of the seven experiments we manipulated several independent variables that should have directly affected the amount of noticing. Although the size of the noticing effect was sometimes substantial, it was never consistently statistically significant. (Later we report several analyses based on individual differences, but they do not change the overall picture.) Although the direction of the noticing effect is often that predicted by the noticing hypothesis (positive), we believe that this should be discounted for several reasons. First, the effect is rarely statistically significant, and it is sometimes significantly reversed. Second, as we will describe more fully in the discussion of Experiment 6, there is reason to believe that many of the small positive effects in Table 2 are due to a subtle confound: the layouts in the not-notice condition tend to be a bit more complicated than the layouts in the notice condition, leading to longer responding in the not-notice condition. Third, an analysis of variance of the Loose z-scores using Experiment as an independent variable indicated that the variation in the size of the notice effect across experiments is not reliable, $F(6, 371) = 1.14$, $MSE = .32$. Fourth, although statistical power is modest in any individual experiment, the power over the set of experiments is very high. Based on the Loose z-score analysis, the power to detect a small effect of 0.2 standard deviations is .97. The power to detect an effect as small as .14 standard deviations is .80. In fact, however, over all of the data (when the number of items is 3 or 4), the mean Loose z-score noticing effect is .04, the standard error is .03, and $t(387) = 1.25$ is not significant. Thus, if there is a noticing effect, it is likely to be very small. All but the most dedicated readers may wish to skip ahead to the General Discussion and avoid the experimental details.

Experiment 1

The first experiment was designed to test the hypothesis that readers would use a mental

model to notice relations not mentioned in a text. The notice condition texts described arrangements in which the target item and the most recently mentioned item ended up adjacent to one another in the spatial layout, although their relation was not described explicitly in the text. In the not-notice condition the target and most recently mentioned item were not adjacent in the spatial layout.

Two manipulations were used to encourage subjects to form a spatial model of the layouts. First, half of the subjects read the texts accompanied by pictures that illustrated part of the spatial layout. The reasoning was that the pictures might serve as a frame to guide subjects' spatial representations (Glenberg & Langston, 1992). Stronger noticing effects were expected in the with-picture condition than in the no-picture condition. Also, all subjects performed a diagram verification task. After reading each text, subjects were presented with a diagram of the layout of the items, and their task was to verify whether or not the diagram portrayed the same layout as described in the text. This task was used to encourage subjects to attend to the spatial aspects of the texts.

Glenberg and Langston (1992) assumed that spatial mental models are constructed using the spatial component of working memory. To test this assumption, an additional manipulation was included: size of the putative mental model at the time of the probe. For half of the texts, subjects were probed after the fourth item was added to the layout (but the text then continued with a description of the other two items). For the other half, subjects were probed after the sixth item. We reasoned that subjects are likely to have enough capacity to represent four items, but that six items would be likely to exceed this capacity. So, we expected to find a larger noticing effect in the four-item condition (because the target item should still be represented in the model when the probe is presented) than in the six-item condition.

Method

Subjects. Eighty subjects participated in the experiment in partial fulfillment of a course research requirement.

Materials. Subjects read a total of 51 texts, 48 experimental texts and three practice texts. Of the 48 experimental texts, 32 were critical texts (eight in each of the four conditions) and 16 were filler texts.

Each text described an arrangement of six items. A sample text is presented in Table 1. Associated with each text was a seventh item that could plausibly be a part of the arrangement in the text, and this item was used when the correct answer to the speeded recognition probe was "no, this item was not mentioned in the text."

For with-picture subjects, the texts were accompanied by a picture depicting the spatial arrangement of the first three items in the text (the entire arrangement was never illustrated). For the sentence describing the first item the picture was a box surrounding the name of the first item. When the sentences describing the second and third items were presented, the picture was updated by adding boxes appropriately located relative to the first box. The picture remained on the screen, with no further changes, while additional items were described. The picture was not visible when the probe appeared. Thus, although the picture could aid memory, in two ways it forces noticing to be based on a cognitive model. First, the critical most recently mentioned item (fourth or sixth item) was not displayed in the picture. Second, the picture was not visible when the probe appeared.

Subjects also verified a diagram after each text. The diagram was composed of six boxes arranged on the screen with the names of the six items written in the boxes. Half of the time these

diagrams corresponded to the layout described by the text. When the diagram did not match the description in the text, the same spatial layout was used, but two items were chosen at random and their locations in the diagram were switched.

The 48 experimental texts were randomly ordered to begin the experiment (the same three texts were always used for practice). The texts were then randomly assigned to one of the conditions. Conditions were assigned such that in each block of six texts there was one text from each of the critical conditions plus two filler texts. These six conditions were randomly ordered in each block, and there were a total of eight blocks.

For filler texts, the probe was either the name of a seventh item (which did not appear in the text), or it was the name of some item in the text other than the first. Approximately two thirds of the filler texts used the seventh item as the probe (because all of the critical texts required a "yes" response, the majority of the fillers required a "no" response to provide some balance). The remaining one third of the filler texts used an item other than the first as the probe. Probes in the filler texts could appear after any sentence except those describing the fourth and sixth items (to break up the pattern of always probing after these sentences).

Design. The first independent variable, picture presence (picture or no picture), was manipulated between subjects. The second independent variable, size (of the putative mental model at the time of the probe), was manipulated within subjects. For half of the critical texts the fourth item was the most recently mentioned item, and for half of the texts the sixth item was the most recently mentioned item. The third independent variable, noticing condition, was manipulated within subjects. The most recently mentioned item was either adjacent to the probed item in the spatial layout (the notice condition) or not adjacent to the probed item (the not-notice condition). The main dependent variable was response time to the probe. Accuracy for the probe question and the diagram verification question was also recorded.

Procedure. Experimental sentences were presented for a duration specified by the equation: $\text{duration} = 1000 + (\text{ms per letter} * [\text{number of characters in the sentence}])$. The result of the computation was a duration in milliseconds, and sentences were presented for this period of time. Before the experiment, subjects set the ms per letter parameter to a comfortable reading speed using a staircase procedure.

Texts were presented to the subject one sentence at a time. Subjects were instructed to imagine the layout on the table in front of them. At a predetermined point during the presentation of the text, a single word (or very short phrase) probe appeared on the screen with three asterisks on either side of it. Subjects answered "yes" or "no" to the question "was this word or phrase in the text you're currently reading?" Subjects were instructed to respond to these probes as quickly as possible without making errors. Affirmative responses were made with the subject's dominant hand.

For critical texts, the probe word was always the first item in the text. For 4-item texts, the probe appeared immediately after subjects read the sentence describing the placement of the fourth item. These texts continued after the probe with the placement of the fifth and sixth items. For 6-item texts, the probe appeared immediately after subjects read the sentence describing placement of the sixth item. After the entire text was presented, subjects verified whether the arrangement of items in a diagram matched the arrangement in the text.

Results

Four separate analyses were conducted for all of the experiments reported in this paper. Two analyses were performed on median response times. For each subject the observations were

sorted by condition and then a median was computed for each condition. The data were analyzed using both the standard and loose analyses described above. The other two analyses were performed using mean log-trimmed-z scores (Bush, et al., 1993) as described in the introduction. These means were also computed using the standard and loose procedures.

As the results of these four analyses were very similar for all of the experiments, we have elected to report in the text only the results from the loose medians analysis. We chose this analysis because it discarded few observations and because median reaction times are easier to interpret than z-scores. Nonetheless, the critical notice versus not-notice comparisons from the other analyses are reported in Table 2. For all analyses, the significance level was set at 0.05.

The data were analyzed using a three-way analysis of variance. The noticing hypothesis predicts a main effect for noticing: Subjects should respond faster in the notice condition than in the not-notice condition. This main effect was not significant, $F(1,78) = 3.06$, $MSE = 8883$, although subjects did respond faster in the notice condition ($M = 1008$ msec) than in the not-notice condition ($M = 1026$ msec). None of the other main effects or interactions was significant.

A second analysis was conducted using percent correct on the probe as the dependent variable. In this analysis, none of the main effects or interactions was significant. Overall accuracy was 93%.

A final analysis used performance on the diagram verification task as the dependent variable. In this analysis, the main effect for size was significant, $F(1,78) = 28.38$, $MSE = 0.03$. Subjects were more accurate in the four-item condition ($M = 78\%$) than in the six-item condition ($M = 68\%$). The interaction between picture and notice was marginally significant ($p < .06$), $F(1,78) = 3.82$, $MSE = 0.02$. No-picture subjects responded approximately equally accurately in the notice and not-notice conditions, whereas with-picture subjects responded more accurately in the notice condition than in the not-notice condition.

Discussion

The results were not consistent with the noticing hypothesis. Subjects did not respond to the probe significantly faster in the notice condition than in the not-notice condition. The failure of the hypothesis cannot be attributed to subjects not encoding spatial information: Performance on the diagram verification task was well above chance. Also, the standard analysis is based only on those texts in which subjects were successful in the diagram verification task.

Nonetheless, there was a hint (given the direction of the difference between notice and not-notice) that a noticing effect was present. One possibility is that noticing does take place, but that the distance manipulation in the first experiment was too weak. It might be that in both conditions, notice and not-notice, the two objects were either too close together or too far apart for a noticing effect to be detected. The plan for Experiment 2 was to refine the distance manipulation to better cover the range of distances between the most recently mentioned item and the target item.

Experiment 2

The notice variable (we will refer to it as the distance variable in this experiment) was expanded to include five distances. The first distance was equivalent to the notice condition in the previous experiment, in that the items in the layout were all immediately adjacent to one another, and the most recently mentioned item ended up adjacent to the target item.

The second distance was the notice-far condition. For texts in this condition, there was a gap in the layout. This gap was created by describing an item as being placed "way [direction]"

with respect to some other item (e.g. "way left of"). An example is presented in Table 1. In this condition, the most recently mentioned item was still (ostensibly) adjacent to the target item, but there was a gap somewhere else in the arrangement. Depending on how subjects treated the instruction to place an item "way" away from another item, the distance between the most recently mentioned item and the probed item could be identical to that in the notice condition, or, if the "way" relation was interpreted as signifying an ambiguous distance, the most recently mentioned item could be close to, but not adjacent to, the probed item.

The third distance was the far condition. Texts in this condition described a layout with a gap, but in this case the gap separated the most recently mentioned item and the target item. For this condition, the target item and the most recently mentioned item were not adjacent (in the layout), but there were no intervening items between them. The fourth distance was equivalent to the not-notice condition of Experiment 1: The target and most recently mentioned items were not adjacent, and there were intervening items between them. The fifth distance was the not-notice-far condition. For this condition, the most recently mentioned item was not adjacent to the target item, there were intervening items between them, and there was a gap in the arrangement.

We were also a bit worried that the picture condition used in Experiment 1 may have been misleading or confusing. That is, the pictures used boxes surrounding the names of the items, rather than analog representations of the items. Perhaps the subjects were confusing the boxes with the items themselves. To try to alleviate this problem, we added a third picture condition that used unlabeled boxes. We hoped that subjects would treat the boxes as simple spatial markers, rather than the objects themselves.

Method

Subjects. Ninety subjects participated in partial fulfillment of a course research requirement. Data from one subject were lost due to a computer error.

Materials and Procedure. There were 55 texts, 40 critical texts (4 in each condition), 12 filler texts, and three practice texts. Each of the critical texts was presented twice, so that there were a total of eight observations in each condition. The presentation patterns used in the second half were the patterns used in the first half rotated 90° clockwise. An unlabeled-picture condition was added. This was the same as the picture condition in Experiment 1, but the names of the items were not written in the boxes.

Design. The first independent variable, picture condition (picture, no picture, or unlabeled picture), was manipulated between subjects. The second independent variable, distance (notice, notice-far, far, not-notice, and not-notice-far), was manipulated within subjects. The third independent variable, size of the arrangement (four items or six items) was manipulated within subjects. The main dependent variable was response time to the probe. Accuracy for the probe question and the diagram verification question was also recorded.

Results

The data were analyzed using a three-way analysis of variance. The noticing hypothesis predicted a main effect for distance. This main effect was not significant, $F(4,328) = 0.90$, $MSE = 9685$. The means for the five distance conditions were: notice $M = 976$ msec, notice-far $M = 992$ msec, far $M = 976$ msec, not-notice $M = 984$ msec, and not-notice-far $M = 978$ msec. A second analysis was conducted with just the 4-item-notice, 6-item-notice, 4-item-not-notice, and 6-item-not-notice conditions to check the replication of the results of Experiment 1. In this analysis the main effect for notice was not significant, $F(1,82) = 0.52$, $MSE = 10626$, although subjects did respond a bit faster in the notice condition ($M = 976$ msec) than in the not-notice

condition ($M = 984$ msec).

The main effect for picture was significant, $F(2,82) = 5.36$, $MSE = 392815$. Subjects responded faster in the with-picture ($M = 912$ msec) and unlabeled-picture ($M = 954$ msec) conditions than in the no-picture condition ($M = 1078$ msec). The main effect for size was significant, $F(1,82) = 13.55$, $MSE = 17056$. Subjects responded faster when the probe occurred after four items ($M = 965$ msec) than six items ($M = 998$ msec). None of the interactions was significant.

A second analysis was conducted using percent correct on the probe as the dependent measure. In this analysis the main effect for size was significant, $F(1,82) = 9.00$, $MSE = 0.013$. Subjects responded more accurately in the 4-item condition ($M = 90\%$) than in the 6-item condition ($M = 88\%$). The interaction between size and distance was also significant, $F(4,328) = 2.41$, $MSE = 0.010$. For the 4-item conditions subjects were approximately equally accurate at all five distances, but in the 6-item conditions subjects were more accurate in the far, not-notice, and not-notice-far conditions than in the notice and notice-far conditions. None of the other main effects or interactions was significant.

A final analysis used performance on the diagram verification task as the dependent variable. In this analysis, none of the main effects or interactions was significant. Overall accuracy on the diagram verification task was 65%.

Discussion

The results of Experiment 2 fail to support the noticing hypothesis. There was no consistent pattern of differences among the five distance conditions. Furthermore, close examination of Table 2 reveals that there were no distance effects for any of the comparisons between the various pairs of distance conditions, even with the more powerful log-trimmed-z transformation. Perhaps limitations on working memory capacity were making it difficult for subjects to form mental models.

Experiment 3

By hypothesis, spatial mental models are constructed using the visual/spatial sketchpad of working memory (Baddeley, 1986). Construction of the model may be interfered with, however, by the necessity to guide the eyes over the text (and picture). To test this idea, an auditory text presentation condition was added. If noticing takes place in the visuo-spatial medium of working memory, then reducing the interference caused by eye-movements should enhance the noticing effect.

Method

Subjects. Sixty-eight subjects participated in the experiment in partial fulfillment of a course research requirement.

Materials and Procedure. There were 50 texts, 48 experimental texts and two practice texts. Each text described an arrangement of four items. The texts were digitized one sentence at a time using a MacRecorder (in the monaural mode, sampling at a rate of 11 MHz). The texts were read by a female assistant in a normal speaking tone and at a rate approximating conversational speech. Two versions were digitized for the critical texts so that they could be presented in either the notice or not-notice condition.

Texts were presented to the subjects one sentence at a time. In the auditory condition, the

texts were played through the Macintosh's internal speaker. The presentation rate in the visual condition was set to approximate the presentation rate in the auditory condition. For both conditions, probe words appeared on the screen.

Design. The first independent variable, picture presence (picture or no picture), was manipulated between subjects. The second independent variable, modality (reading or listening), was manipulated between subjects. The third independent variable, noticing condition, was manipulated within subjects. In the notice condition item four was adjacent to item one (the target item), and in the not-notice condition item four was not adjacent to item one. The main dependent variable was response time to the probe. Accuracy for the probe and diagram verification questions was also recorded.

Results

The data were analyzed using a three-way analysis of variance. According to the noticing hypothesis, subjects should respond faster in the notice condition than in the not-notice condition. This main effect was not significant, $F(1,64) = 0.55$, $MSE = 14045$. However, subjects did respond a bit faster in the notice condition ($M = 983$ msec) than in the not-notice condition ($M = 998$ msec). None of the other main effects or interactions was significant.

A second analysis was conducted using percent correct on the probe as the dependent variable. In this analysis, none of the main effects or interactions was significant. Overall accuracy was 93%.

A final analysis used performance on the diagram verification task as the dependent measure. In this analysis, the main effect for picture was significant, $F(1,64) = 10.92$, $MSE = 0.03$. Subjects were more accurate in the with-picture condition ($M = 85\%$) than in the no-picture condition ($M = 75\%$). The main effect for modality was also significant, $F(1,64) = 25.58$, $MSE = 0.03$. Subjects were more accurate in the auditory condition ($M = 88\%$) than in the visual condition ($M = 72\%$). The main effect for notice was also significant, $F(1,64) = 7.08$, $MSE = 0.01$. Subjects were more accurate in the not-notice condition ($M = 82\%$) than in the notice condition ($M = 77\%$). The interaction between picture and modality was marginally significant, $F(1,64) = 3.81$, $MSE = 0.03$, $p = .06$. Subjects in the with- and no-picture conditions were more accurate with auditory presentation than visual presentation, but the effect was more pronounced for subjects in the no-picture condition. None of the other interactions was significant.

Discussion

The results from Experiment 3 provide little support for the noticing hypothesis. Subjects showed about the same noticing effect in the auditory and visual conditions, and the overall difference was very small. Performance on the diagram verification task indicates that the picture and modality manipulations had the desired effect. That is, subjects represented the spatial layout of the items more accurately in the conditions designed to reduce interference with spatial working memory, but there was still no noticing effect in those conditions.

Experiment 4

Anonymous reviewers noted a number of potential methodological problems with Experiments 1-3. The reviewers' principle objections were as follows. 1) Item 1 was always the critical probe item. This may engender guessing strategies and it may encourage special effort to keep the item activated, thus overriding any effect of noticing. 2) The correct answer to the recognition probe is more frequently "yes" than "no," suggesting that subjects might adopt a strategy to respond "yes" most of the time. 3) The evidence that the item recognition task is

tapping a spatial representation is not as convincing as it might be. 4) The evidence that subjects are forming an appropriate representation of the spatial layouts of the texts as indexed by diagram verification is convincingly above chance for only some of the conditions.

Insert Table 3 About Here

Experiment 4 was designed to address these objections. All of the texts described arrangements of three items. An example is given in Table 3. To correct the first problem, half of the notice and not-notice probes were item one and half of the notice and not-notice probes were item two. Item three was also used as the probe item with the same frequency as items one and two to eliminate any potential focusing strategy that might be employed by the subjects. However, data from the item three probes did not enter into the analyses. To correct the second problem, there were an equal number of "yes" and "no" probes.

To correct the third problem, we used a set of "spatial" probes² to demonstrate that the item recognition probes are tapping into a spatial representation. From the subject's point of view, these spatial probes were identical to the notice and not-notice probes; that is, the spatial probes were item recognition probes. Half of the spatial probes involved an item at an end of the dimension and half of the spatial probes involved an item in the middle of the dimension. Items one and two were used as the spatial probes equally often. If there is a difference between response times to items from the end of the dimension and items in the middle, this would indicate that the item recognition task is sensitive to spatial location. If there is an effect of spatial location in the absence of a noticing effect, this would indicate that the lack of a noticing effect is not due to insensitivity of the item recognition probe to spatial representations.

To correct the fourth problem, we replaced the diagram verification task with a symbolic distance task. The symbolic distance effect is that when items are arranged along a dimension (the dimension is "comfort" for the example in Table 3), judgments about pairs of items that are near on the dimension take longer than judgments about pairs of items that are far apart on the dimension (Potts, 1972; Moyer, 1973). For the text in Table 3, subjects should verify "planes seem more comfortable than trains" faster than "planes seem more comfortable than cars" since planes and trains are farther apart on the dimension of comfort. If subjects in this experiment are forming an appropriate representation of the spatial arrangement of the items in the text, then we ought to get a symbolic distance effect. Note that because the order in which the items are described in the text is different from their order on the dimension, correct performance on the symbolic distance task requires forming the appropriate spatial ordering.

The design of the experiment is not ideal for examining the symbolic distance effect. The problem is that there is a confounding of "distance" between items on the dimension and whether the items are on the ends of the dimension. This confounding is not a problem here, however, because whether the symbolic distance effect is due to end items or distance is irrelevant. In either case, observing the effect is sufficient to conclude that readers were constructing the appropriate spatial representation because both "end item" and "distance" are defined in terms of the constructed representation, not the order of occurrence of the items in the text.

Method

Subjects. Twenty-six subjects participated in the experiment. Two subjects failed to complete the experiment, and data from two subjects were lost due to computer failure.

Materials. There were 98 texts, 96 experimental texts and two practice texts. Of the 96

experimental texts, eight had notice probes, eight had not-notice probes, eight had spatial-end probes, eight had spatial-middle probes, 16 had item three probes, and 48 had "no" probes.

Each text described an arrangement of three items. Item one and item two probes were either notice, not-notice, spatial-end or spatial-middle probes. There was an equal number of "no" and item three probes. There are several features to note about these texts. First, the texts describe a pattern in the front/back dimension a) because this dimension appears to be easier than the left/right dimension to represent and access (Bryant, Tversky, & Franklin, 1992), and b) this dimension should reduce stimulus-response compatibility problems when using left and right hands to respond "yes" and "no." Second, the relation between the last mentioned item (item three) and the target is always implicit, rather than given directly in the text. Thus we can ask if accessibility of the target varies with its spatial proximity to the last mentioned item. Third, the locution "way in front of" or "way behind" was used in every text to describe the relation between the first and second items. This was done to leave room (in a spatial layout) for the possible insertion of the object named by the third item. Because the results from Experiment 2 indicated that near versus far in a spatial model (with no intervening items) had little effect on noticing, we were not concerned by the introduction of this locution.

After each text, subjects responded to six true/false symbolic distance questions. Each question was a statement about the relative positions of a pair of items from the text on the dimension. Sample questions are presented in Table 3. The six questions contained all possible combinations of the three items in the text, and the order of the questions was randomized independently for each text for each subject. For half of the questions the correct response was "true" and for half of the questions the correct response was "false".

Design. All independent variables were manipulated within subjects. The independent variables associated with the item-recognition probes were noticing condition (notice, not-notice), location (end, middle), and item (one, two). The independent variable associated with the symbolic distance questions was distance (near, far). The dependent measures were response time and accuracy.

Procedure. Subjects were presented with instructions via computer. Each text was presented one sentence at a time, and the presentation time for each sentence was determined by the equation $[\text{time} = (\text{Length}(\text{sentence}) * 33) + 2000]$ msec. After presentation of the sentence describing the location of the third item, subjects responded to an item-recognition probe. Subjects always responded "yes/true" with their dominant hand. Incorrect responses were followed by a beep. After responding to the probe, subjects responded to the six symbolic distance sentences.

Results

Three separate analyses were performed. First, the data from the spatial probes were analyzed using a two-way ANOVA with location and item as the factors. According to the spatial hypothesis, we would expect to see a main effect for location. This main effect was significant, $F(1,21) = 6.22$, $MSE = 13408$. Subjects responded faster to the middle item ($M = 1012$ msec) than to the end item ($M = 1074$ msec). The main effect for item and the interaction between location and item were not significant. These data indicate that the item recognition task used in the experiment was sensitive to a spatial representation. A second ANOVA was conducted using percent correct for the spatial probes as the dependent measure to check for a speed accuracy trade-off. In this analysis, none of the effects was significant. Overall accuracy was 96%.

If subjects are forming a lasting representation of the spatial arrangement, then we ought to see a symbolic distance effect. In particular, subjects ought to respond faster to sentences

containing pairs of items that are far on the dimension than to sentences containing pairs of items that are near on the dimension. This effect was significant, $t(21) = 6.17$, $SEM = 44.93$. The mean for far sentences was 1294 msec, and the mean for near sentences was 1572 msec. These data indicate that subjects are forming an appropriate representation of the arrangements described in the texts. An analysis of accuracy data from the symbolic distance questions also showed an effect for distance, $t(21) = -3.91$, $SEM = 0.011$. Subjects were more accurate in the far condition ($M = 89\%$) than in the near condition ($M = 85\%$).

The third analysis investigated the noticing effect. The data were analyzed using a two-way ANOVA³ with notice condition and item (the target item was item 1 or 2) as the factors. If the noticing hypothesis were correct, we would expect subjects to respond faster to notice probes than to not-notice probes. This main effect was not significant, $F(1,21) = 0.28$, $MSE = 18782$, although subjects did respond a bit faster in the notice condition ($M = 1074$ msec) than in the not-notice condition ($M = 1090$ msec). The main effect for item and the interaction were not significant. An analysis of the accuracy data showed no main effects or interactions. Overall accuracy was 95%.

Discussion

All items were used as the probe item equally often and there were equal numbers of "yes" and "no" responses, ruling out the possibility that subjects were using simple guessing strategies. Data from the symbolic distance task indicate that subjects were clearly forming an accurate representation of the arrangements described in the texts. The data from the spatial probes indicate that the item recognition task was sensitive to spatial arrangement. In spite of this, the results of Experiment 4 fail to support the noticing hypothesis.

One component of these data seems a bit unusual. For the spatial probes, responding was faster for items in the middle of the ordering than for items on the ends. The expectation of an advantage for end items arises, in part, from work on the symbolic distance effect: Symbolic distance comparisons involving end items are fast. However, the spatial probes did not involve comparisons, only access. To our knowledge, there is no data demonstrating faster access to end items using a recognition probe. Faster responding to middle items is consistent with a spatial representation if it is assumed that subjects are focusing on the middle items in a spatial representation (perhaps looking at them with the mind's eye) in an attempt to keep all items equally available.

Perhaps our inability to find a noticing effect reflects the use of an insensitive dependent variable. In Experiments 5-7 we used as a dependent variable the time to read a sentence containing an anaphoric reference to the target item.

Experiment 5

The dependent variable was the time to read an about sentence. This sentence always provided a new fact about the target item. An example is given in Table 4. The about sentence was presented immediately following the description of the fourth item which was either adjacent to the target item in the spatial layout (notice condition) or separated from the target item (not-notice condition). We hypothesized that when constructing the model subjects would be more likely to access the target item in the notice condition than in the not-notice condition, and that this noticing should speed reading of the about sentence in the notice condition.

Insert Table 4 About Here

Following the about sentence was a move sentence. It described how the protagonist moved one of the previously located items (the target item for the critical texts) to a new location. The move sentence was included to ensure that subjects were building a mental model that was capable of being updated, rather than a static representation.

Method

Subjects. Fifty subjects participated in the experiment in partial fulfillment of a course research requirement.

Materials. Thirty-five texts, 32 experimental texts and three practice texts, were used in the experiment. The experimental texts were each used twice and were randomly reassigned to conditions for the second pass. This resulted in 16 notice texts, 16 not-notice texts, and 32 filler texts for each subject. Each text described an arrangement of four items. After the description of the location of the fourth item, an about sentence was presented. This sentence described something about the target item. This was followed by a move sentence describing the movement of an item to another location.

Design. A single factor within-subjects design was used. The independent variable was noticing condition. In the notice condition, item four was adjacent to item one (the target item) in the spatial layout, and in the not-notice condition item four was not adjacent to item one. The main dependent variable was reading time for the about sentence. Accuracy in responding to a diagram verification task was also recorded.

Procedure. Subjects began the experiment with three practice texts. If the subject answered incorrectly on the diagram verification for all three of the practice texts they were prompted to alert the experimenter, who made sure that they understood the procedure. Texts were presented to the subjects one sentence at a time. The presentation rate was self-paced. After the sentence describing the location of the fourth item, the about sentence was presented, and reading time was collected.

After reading the text, subjects answered the diagram verification question. The diagram verification task was in a multiple choice format. Four layouts for the items in the text were displayed, and the subject selected the one that correctly represented the layout after the move sentence. Subjects used the keyboard to type their answers to the comprehension question.

Results

The difference between the notice and not-notice conditions was not significant, $t(49) = 0.03$, $SEM = 55.4$. Subjects' median reading time in the notice condition ($M = 1538$) was virtually the same as in the not-notice condition ($M = 1539$).

A second analysis was performed using percent correct on the diagram selection task as the dependent variable. The difference between the percent correct for the two conditions was significant, $t(49) = 4.94$, $SEM = 0.03$. Subjects were more accurate in the notice condition ($M = .75$) than in the not-notice condition ($M = .59$).

Might the degree of noticing be related to individual differences? In several of the earlier experiments we investigated whether noticing was correlated with individual differences in comprehension or individual differences in spatial memory capacity. We did not report this work for the earlier experiments because we did not find any significant relations that replicated across experiments. We take up the question again here, in part because it provides a major motivation for the next experiment. We reasoned as follows. Noticing presupposes that people are building

mental models and comprehending the texts. In fact, if individuals are not attempting to comprehend the spatial layouts, it is unlikely that they would show a noticing effect. Thus, we predicted a positive correlation between size of the noticing effect and performance on the diagram selection task. The correlation was not significant, however, $r(48) = 0.14$.

Discussion

The results from Experiment 5 provide little support for the noticing hypothesis. Nonetheless, we were encouraged (perhaps perversely) by the correlational analysis. Although the correlation from the loose medians analysis was not significant, other versions of the correlation flirted with statistical significance. Thus, we thought it worth another try.

The correlation might have a source having little to do with noticing, however. The not-notice conditions have arrangements that are, in some ways, more complex than the arrangements used in the notice conditions. For example, most of the patterns in the notice condition conform to a square or rectangular layout of objects, whereas many of the patterns in the not-notice condition are more akin to a z-shape. Perhaps the not-notice patterns are more difficult to construct or hold in memory. This difference in pattern complexity could explain the trend to have slower reaction times and slower reading times in the not-notice conditions (see Table 2). Differences in pattern complexity might also give rise to the correlation between the size of the noticing effect and performance on the diagram selection task. Consider first a subject who is working hard to comprehend. This subject will do well on the diagram selection task. Also, because the subject is representing the spatial layouts while reading, the subject will read the complex not-notice texts more slowly than the less complex notice texts. That is, the subject will produce a noticing effect by virtue of complexity affecting reading time, not noticing. Now consider a subject who is not working diligently to comprehend. This subject will perform poorly on the diagram selection task. In addition, because this subject is not constructing the spatial layouts, the subject will not be adversely affected by the complex not-notice texts. Considering the patterns produced by these two types of subjects, a positive correlation between performance on the diagram selection task and the size of the noticing effect might be due to a relation between comprehension and susceptibility to the more difficult not-notice arrangements, rather than a relation between comprehension and noticing on the basis of distance between elements in a spatial mental model.

Experiment 6

The design of Experiment 6 allowed us to discriminate between the two explanations for the correlation between noticing and performance on the diagram selection task. In the target-item-one condition, subjects read exactly the same texts as did the subjects in Experiment 5. If the correlation between the size of the noticing effect and performance on the diagram selection task is due to noticing, then subjects in the target-item-one condition should produce a noticing effect and a correlation between the size of the noticing effect and performance on the diagram selection task. In the target-item-two condition, subjects read texts that were modified so that the about sentence described a fact about the second item in the text. When item two is the target, the most recently mentioned item (item 4) is diagonally removed from the target in both the notice and not-notice conditions. Thus, the distance between the most recently mentioned item and the target item is the same in the notice and not-notice conditions. In this case, if distance between elements in a mental model is what produces a noticing effect, there should be no noticing effect in the target-item-two condition. Nonetheless, in the target-item-two condition the not-notice arrangements are more complex than the notice arrangements. Thus, if complexity of arrangements is what produces the "noticing effect" (in quotes because it is an effect of complexity, not noticing), subjects in the target-item-two condition should produce a "noticing effect" and a correlation between the size of the "noticing effect" and performance on the diagram selection task. To summarize, if distance in the mental model is the operative variable, only subjects in the target-item-one condition should

produce a noticing effect and a correlation. On the other hand, if complexity of arrangement is the operative variable, subjects in both conditions should produce the "noticing effect" and the correlation.

Method

Subjects. Seventy-eight subjects participated in the experiment in partial fulfillment of a course research requirement.

Materials and Procedure. Half of the subjects (those in the target-item-one condition) saw texts that were identical to those in Experiment 5. Importantly, the about sentence (for which we collected reading time) referred to item 1. For the other subjects (those in the target-item-two condition) the about sentence pertained to item two.

Design. A mixed-factorial design was used. The between-subjects independent variable was target item. In the target-item-one condition the about sentence referred to item one. In the target-item-two condition the about sentence referred to item two. The within-subjects independent variable was noticing condition. In the notice condition, item four was adjacent to item one in the spatial layout, and in the not-notice condition item four was not adjacent to item one. The main dependent variable was reading time for the about sentence. Accuracy in responding to the diagram selection task was also recorded.

Results

The data were analyzed using a two-way analysis of variance. According to the hypothesis, subjects in the target-item-one condition should read notice texts faster than not-notice texts, but subjects in the target-item-two condition should not show any difference between notice and not-notice texts. Thus, the hypothesis predicts an interaction between noticing condition and target item. This interaction was not significant, $F(1,76) = 0.08$, $MSE = 64909$. The reading times for target-item-one showed the predicted pattern in that the about sentence was read faster in the notice condition ($M = 1615$ msec) than not-notice condition ($M = 1652$ msec). However, approximately the same difference was found in the target-item-two condition: The about sentence was read faster in the notice condition ($M = 1606$ msec) than in the not-notice condition ($M = 1666$ msec). No main effects or interactions were significant (for the main effect of noticing, $F(1,76) = 1.41$, $MSE = 64909$).

A second analysis was performed using percent correct on the diagram selection task as the dependent variable. The main effect for noticing was significant, $F(1,76) = 14.58$, $MSE = 0.02$. Subjects were more accurate in the notice condition ($M = .74$) than in the not-notice condition ($M = .65$). No other main effects or interactions were significant.

Correlations between the size of the noticing effect and performance on the diagram selection task were computed separately for target-item-one and target-item-two conditions. They were $-.11$ and $-.04$, respectively, and neither was statistically significant.

Discussion

The results from Experiment 6 provide little support for the noticing hypothesis. The difference between notice and not-notice median times to read the about sentence were in the predicted direction. But a) the difference was not significant, b) the same magnitude of a difference was found in the target-item-two condition which should not have produced a noticing effect according to the noticing hypothesis, and c) the correlation between size of the noticing effect and performance on the diagram selection task was not significant, and it was not even in the

right direction. As in Experiment 5, there was a significant difference between the notice and not-notice conditions for the diagram verification task. Apparently, the not-notice layouts are a bit more complex than the notice layouts. This difference in complexity may well underlie the small positive noticing effects reported in Table 2.

Experiment 7⁴

In Experiment 7 we manipulated a frame of reference variable (as in Logan, 1995). In all of the previous experiments, the most recently mentioned item is located in reference to another item. Thus, using the example text in Table 5, the box of chicken is located to the left or right of the picnic basket. In other words, to locate the most recently mentioned item, attention must be directed to the other item. Perhaps, then, any noticing is in respect to that other item, the item around which the frame of reference is constructed. In this experiment, we manipulated whether the frame of reference remained on the other item (item 2), or whether it was shifted to the most recently mentioned item (item 3). This manipulation was accomplished by adding to the text a verification sentence between the sentence locating item 3 and the about sentence. The verification sentence restated the relation between item 3 and item 2, but in a manner that could shift the frame of reference. Using the example in Table 4, with the verification sentence "So, the box of chicken is on the picnic basket's left," the frame of reference remains on item 2, the picnic basket. Using the verification sentence, "So, the picnic basket is on the box of chicken's right," the frame of reference is located on item 3, the most recently mentioned item. Both versions of the verification sentence describe the same spatial layout. The slightly odd wording for the verification sentence was chosen to avoid verbatim repetition of the previous sentence in the not-notice condition. The frame-of-reference-modified noticing hypothesis predicts greater noticing when the frame of reference is shifted to item 3 (the most recently mentioned item) than when the frame remains on item 2.

A second issue was in how we encouraged subjects to maintain a spatial layout. In the previous experiments, we tested memory for spatial layout and depended on the subject's motivation to do well on the test. In this experiment we upped the motivational ante. After reading a text, subjects responded to two questions about the spatial layout and two questions about the about sentences. If the subject missed one or more of these questions, the condition was repeated later in the session. Subjects were forewarned of this, and they were given feedback following each text. Of course, we did not want to have some subjects continuing indefinitely, and so (unbeknownst to the subjects) the maximum number of texts that could be repeated was 10.

Method

Subjects. Forty subjects participated in the experiment in partial fulfillment of a course research requirement. Two subjects withdrew from the experiment and the data from a third subject were lost due to computer failure.

Materials. Subjects read, initially, a total of 43 texts, 40 experimental texts (32 critical texts and 8 filler texts) and three practice texts. Each text described an arrangement of three items. Each text also contained a verification sentence, which served to put the frame of reference on the desired item, and two about sentences, one referring to item one, the other referring to item two. The initial about sentence always referred to item 1, and we measured the reading time of this sentence as an index of noticing. The second about sentence was included so that the initial about sentence would not be the last sentence in the passage.

Instead of a diagram selection test, subjects answered two test questions regarding the spatial layout described in the text. The specific questions were randomly selected for each subject and text so that the correct answer was approximately equally often yes or no. Additionally,

subjects responded to two "yes" or "no" questions regarding the about sentences.

Design. Both independent variables were manipulated within subjects. The first independent variable was frame of reference. In the item-two-frame-of-reference condition the verification sentence was worded to place the frame of reference on item two. In the item-three-frame-of-reference condition the verify sentence was worded to place the frame of reference on item three. The other independent variable was noticing condition. In the notice condition, item three was adjacent to item one (the target item) in the spatial layout, and in the not-notice condition item three was not adjacent to item one. The main dependent variable was reading time for the initial about sentence. Accuracy in responding to the four questions was also recorded and used to determine the necessity for repetition of a text in a particular condition.

Procedure. Texts were presented one sentence at a time. Presentation rate was self-paced. Subjects were required to correctly answer all questions on the three practice texts before going on to the experimental texts. For each text, the subjects answered two yes/no questions regarding the layout of the items described in the text, as well as a yes/no question for each of the two about sentences. If any of these four questions was missed, the subject had to read an additional text in that condition following completion of the initial 40 experimental texts. The particular text reread was selected at random from the 40 experimental texts, but was presented in the same condition as the text for which the subject answered a question incorrectly.

Results

If subjects missed any of the test questions, the data from that text were dropped and the condition repeated after the initial pass through the 40 texts. Subjects reread an average of 3.27 (2.27) notice texts, 2.38 (2.24) not-notice texts, and 2.03 (2.02) filler texts. By the end of the experiment, 30 subjects had contributed eight observations to each of the critical conditions, and 7 subjects had fewer than eight observations in each critical condition either because they had missed questions on more than 10 texts in the initial pass, or because they missed questions during the second pass. We used the data from all 37 subjects to maximize statistical power.

The times to read the initial about sentences were analyzed using a two-way analysis of variance. The frame-of-reference-modified noticing hypothesis predicts a) a main effect of noticing condition and b) an interaction between notice and frame of reference. The main effect was not significant $F(1,36) = .62$, $MSE = 107867$, nor was the interaction, $F(1,36) = 0.02$, $MSE = 110841$. In the item-three-frame-of-reference condition, subjects read the initial about sentence faster for notice texts ($M = 1762$ msec) than for not-notice texts ($M = 1797$ msec). The same was true for the item-two-frame-of-reference condition, notice text $M = 1762$ msec, not-notice text $M = 1812$ msec.

Discussion

Although there is a hint of a noticing effect in the z-score analysis (see Table 2), because it is of borderline significance, because we have done so many analyses, and because the interaction with frame of reference was not close to being significant, we believe that the apparent effect of noticing is not real.

General Discussion

The empirical conclusion is inescapable: Noticing rarely occurs. To say it differently, readers do not seem to infer relations in a mental model on the basis of spatial contiguity alone. The caveat, "alone" is important. One can imagine situations in which a reader might well infer unstated spatial relations using past experience (e.g., the description of a familiar room) or when

the spatial relations are particularly important. In any event, we have demonstrated that simply forming a cognitive representation of stated spatial relations does not engender noticing of additional spatial relations.

It is also clear from our data as well as the data from many other laboratories that readers encode spatial relations and that they are functional. Notice the good performance on the various types of spatial comprehension questions as well as the symbolic distance effect reported for Experiment 4. It appears, however, that the sorts of spatial relations encoded from text are all explicitly stated in the text, are depicted in an accompanying picture, or are based on repeated presentations of the text and task demands to form a detailed spatial representation (e.g., Denis & Cocude, 1992; Glenberg & Langston, 1992; Rinck, Hähnel, Bower, & Glowalla, 1995; Wagener & Wender, 1985). There is little evidence that readers infer spatial relations on the basis of contiguity in a spatial mental model; that is, there is little evidence for noticing.

The prediction that noticing would occur is derived from a system of assumptions. Which of those assumptions is incorrect? For example, it may well be that noticing does not occur, but that mental models are, nonetheless, built in a Euclidean-like medium. This seems unlikely for several reasons. First, without a mechanism like noticing, that is, without a mechanism for making spatial distances functional in and of themselves, the idea of a Euclidean-like representation loses much force. If the spatial distances are not functional, then operations on the Euclidean-like representation must invoke a different process that probes the representation, a mind's eye, if you will, or in more pejorative terms, a homunculus. Postulating such a process then requires other levels of explanation (does the mind's eye have an eye?), leading to a type of infinite regress.

Second, there is now a body of data questioning the claim that the mental representation of spatial relations is Euclidean. McNamara (1986) demonstrated that after learning locations of objects, objects near to one another primed each other more effectively than those farther away. However, objects in a single bounded region led to greater priming than equally distant objects in two regions. In other words, spatial priming was not a simple function of distance; region also affected the degree of relation. McNamara, Hardy, and Hirtle (1989) were able to show that subjective regions (as opposed to those marked explicitly on a map) also affected spatial priming. The conclusion is that spatial information is not represented in a system that is purely Euclidean, but that hierarchical or topological relations (e.g., regions) also affect performance.

Tversky and her colleagues (e.g., Bryant, et al., 1992; Franklin & Tversky, 1990) have proposed that spatial descriptions are retrieved (and perhaps encoded) using a "spatial framework" that is decidedly not Euclidean. Their basic findings come from a paradigm in which readers first read about and memorize a three-dimensional spatial layout of objects. The reader is then instructed to imagine facing one of the objects (or the protagonist in the text is described as facing one of the objects). Finally, the reader is instructed to retrieve the names of objects in various directions such as up/down, front/back, and left/right. The basic finding is that retrieval times are a function of the direction, with the fastest times being for up/down and the slowest for left/right. Once again, the conclusion is that the representation of spatial layout (or the retrieval of objects from the representation) does not reflect an unstructured space.

There is also reason to question the Euclidean interpretation of the distance effects observed in the Morrow and Bower paradigm (e.g., Morrow et al., 1989). To review, in Rinck and Bower (1995) and Rinck et al. (1995, from which the following examples are taken), subjects memorized the spatial layout of a building and then read about the movements of a protagonist through the building. A motion sentence such as "Then he walked from the storage room into the lounge" described the movement of the protagonist from a source room (storage room) to a goal room (lounge), and the sentence implied a path that traversed an unnamed path room (in this case, a repair shop). After reading the motion sentence and a motivating sentence that provided a context

for the target sentence, the subjects read a target sentence that referred to an object in the source, path, or goal rooms. An example of a target sentence is, "He decided that the cart should not be so dirty tomorrow." Time to read the target sentence was a monotonically increasing function of distance of the room containing the object from the protagonist's location in the the goal room. This result is consistent with the claim that spatial (and perhaps Euclidean) distance from the protagonist is a critical factor in comprehending reference to objects. However, Rinck et al. (1995) report data that strongly question the idea of a Euclidean representation affecting text comprehension. In these experiments, the memorized layout included "path" rooms that were either divided (so that there were two path rooms implied by the motion sentence) or undivided. The divided and undivided path rooms contained the same objects and occupied the same Euclidean space. Rinck et al. tested if the time to read a sentence referring to an object in the path room was sensitive to distance of the object (in the path room) from the protagonist. In fact, when the path room was divided, distance affected reading time so that sentences about objects in the path room far from the protagonist took longer to read than sentences about objects in the path room closer to the protagonist. This effect could either be due to Euclidean distance or a category effect based on the number of rooms. The critical conditions involved the undivided path room containing objects that were literally as far from the the protagonist as when the path room was divided. In this case, Euclidean distance played no role in reading time. Apparently, number of rooms, rather than Euclidean distance, produces the effect. Interestingly, Rinck et al. did find evidence that spatial representations can incorporate Euclidean components. When the subjects were asked to judge the relative distance of objects from the protagonist, the speed of the judgments was inversely related to Euclidean distance between the objects. This result corresponds to the finding in Experiment 4 that subjects performance on the symbolic distance task seemed to reflect distance. Overall, the results from the Rinck et al. experiments imply that spatial information derived from a memorized layout may be represented with Euclidean properties, but that those properties do not affect text comprehension.

Much of the data appear to be covered by a relatively simple generalization. Cognitive representations derived from perception, pictures, or many repetitions of a text (see Denis & Cocude, 1992), may well have Euclidean components, but the Euclidean nature of those components do not seem to affect text comprehension. Furthermore, on the basis of the experiments reported here, reading in the absence of pictures or multiple passes at the text is unlikely to result in a representation that is Euclidean.

We began this investigation with questions about the nature of mental models. Can we say anything more positive about that nature? Again, one conclusion is clear: Although the idea of an unstructured, Euclidean spatial layout is attractive for its simplicity, it appears to be wrong. An alternative is to conclude, along with van Dijk and Kintsch (1983), that mental models are intrinsically propositional. There are strong grounds to question this alternative, however (see Barsalou, 1993; Glenberg, in press; and Lakoff, 1987). In brief, propositions are composed of abstract symbols that are difficult or impossible to associate with external referents. Thus, propositional systems are inherently syntactic and represent only relations to other propositions, not relations to the world. Mental models, on the other hand, are quintessentially semantic; they encode our understanding of particular situations in the world, not just relations among abstract symbols.

After discounting Euclidean spatial representations and propositional representations, there are still many possibilities, although none are currently as well-developed as the propositional account. Several of these are related to the concept of embodiment (Johnson, 1987; Lakoff, 1987; see Rinck et al., 1995, for other possibilities). In contrast to propositional representations, embodied representations are not constructed from abstract symbols. Instead, the basic elements of the system are shaped by how the body interacts with the environment. Thus, Johnson (1987) describes how the physical nature of our bodies leads to consistent, structured, experiences of in-

out, such as putting food into our bodies and moving into and out of rooms. This sort of basic bodily experience is proposed to underlie our understanding of concepts such as containment, as well as our understanding of more metaphorical uses of spatial terms such as to wake out of a deep sleep. In contrast to representations built out of Euclidean space, embodied representations are highly structured. For example, Lakoff (1987) describes the structure of the container schema as having an inside, an outside, and a boundary. Importantly, the schema specifies how being in a container carries many implications (e.g. that what is in the container is under the control of the container). Embodied accounts of meaning are being developed by Lakoff and Johnson (e.g., Lakoff, 1987), Barsalou (1993), and Glenberg (in press).

Consider how such an account might handle the contrast between the results of Glenberg et al. (1987), which suggest a strong contribution of spatial relations to comprehension, and the results of the current experiments suggesting the opposite. When the protagonist, John, puts on his sweatshirt to go jogging, the sweatshirt is spatially close to him. What may be more important, however, is that the sweatshirt is literally attached to John. Attachment is just the sort of relation that can arise from bodily interactions with the world, such as holding hands with one's parents (see Lakoff's discussion of the "Link" schema). One of the consequences of attachment is a dependency, so that that the sweatshirt goes where John goes. When John takes off his sweatshirt, the attachment relation is broken and the dependency no longer applies. Given these functionally different relations, we would suspect that having the sweatshirt on or off would affect comprehension, and it does. In contrast, consider the spatial arrangement described by the sample text in Table 1. Whereas the objects have different Euclidean relations to one another, there is little or no differentiation in regard to functional, topological, or embodied relations. For example, all of the objects in the text in Table 1 can be considered to be in the same container (the aquarium), and thus all are equally likely to interact with each other and equally likely not to interact with objects outside of the aquarium. On this analysis, and consistent with the experiments reported here, because the objects are not differentiated in regard to embodied relations, they are not differentiated in regard to effects on comprehension.

References

- Baddeley, A. D. (1986). *Working Memory*. Oxford: Oxford University Press.
- Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, M. A. Conway, & P. E. M. Morris (Eds.), *Theories of memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bryant, D. J., Tversky, B., & Franklin, N. (1992). Internal and external spatial frameworks for representing described scenes. *Journal of Memory and Language*, 31, 74-98.
- Bush, L. K., Hess, U., & Wolford, G. (1993). Transformations for within-subject designs: A Monte Carlo investigation. *Psychological Bulletin*, 113, 566-579.
- Denis, M., & Cocude, M. (1989). Scanning visual images generated from verbal descriptions. *European Journal of Cognitive Psychology*, 1, 293-307.
- Denis, M., & Cocude, M. (1992). Structural properties of visual images constructed from poorly or well-structured verbal descriptions. *Memory and Cognition*, 20, 497-506.
- Denis, M., Gonçlaves, M-R, & Memmi, D. (in press). Mental scanning of visual images generated from verbal descriptions: Towards a model of image accuracy. *Neuropsychologia*.
- van Dijk, T. A., & Kintsch, W. (1983). *Strategies of discourse comprehension*. New York: Academic Press.
- Franklin, N., & Tversky, B. (1990). Searching imagined environments. *Journal of Experimental Psychology: General*, 119, 63-76.
- Garnham, A. (1992). Minimalism versus constructionism: A false dichotomy in theories of inference in reading. *PSYCOLOQUY*, 3(63), reading-inference-1.1.
- Gentner, D., & Stevens, A. L. (Eds.). (1983). *Mental models*. Hillsdale, NJ: Erlbaum.
- Glenberg, A. M. (in press). What memory is for. *Behavioral and Brain Sciences*.
- Glenberg, A. M., Kruley, P., & Langston, W. E. (1994). Analogical processes in comprehension: Simulation of a mental model. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*. Orlando, FL: Academic Press.
- Glenberg, A. M., & Langston, W. E. (1992). Comprehension of illustrated text: Pictures help to build mental models. *Journal of Memory and Language*, 31, 129-131.
- Glenberg, A. M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Language*, 26, 69-83.
- Graesser, A. C., Singer, M., & Trabasso (1994). Constructing inferences during narrative text comprehension. *Psychological Review*, 101, 371-395.
- Johnson, M. (1987). *The body in the mind*. Chicago, IL: The University of Chicago Press.
- Johnson-Laird, P. N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.

- Johnson-Laird, P. N., Byrne, R. M. S., & Tabossi, P. (1989). Reasoning by model: The case of multiple quantification. *Psychological Review*, 96, 658-673.
- Lakoff, G. (1987). *Women, fire, and dangerous things*. Chicago: University of Chicago Press.
- Logan, G. (1995). Linguistic and conceptual control of visual spatial attention. *Cognitive Psychology*, 28, 103-174.
- Mani, K., & Johnson-Laird, P. N. (1982). The mental representation of spatial descriptions. *Memory & Cognition*, 10, 181-187.
- McNamara, T. P. (1986). Mental representation of spatial relations. *Cognitive Psychology*, 18, 87-121.
- McNamara, T. P., Hardy, J. K., & Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 211-227.
- Morrow, D. G., Bower, G. H., & Greenspan, S. L. (1989). Updating situation models during narrative comprehension. *Journal of Verbal Learning and Verbal Behavior*, 28, 292-312.
- Moyer, R.S. (1973). Comparing objects in memory: Evidence suggesting an internal psychophysics. *Perception and Psychophysics*, 13, 180-184.
- O'Brien, E. J., & Albrecht, J. E. (1992). Comprehension strategies in the development of a mental model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 777-784.
- Payne, S. J. (1993). Memory for mental models of spatial descriptions: An episodic-construction-trace hypothesis. *Memory & Cognition*, 21, 591-603.
- Potts, G.R. (1972). Information processing strategies used in the encoding of linear orderings. *Journal of Verbal Learning and Verbal Behavior*, 11, 727-740.
- Rinck, M., & Bower, G. H. (1995). Anaphora resolution and the focus of attention in situation models. *Journal of Memory and Language*, 34, 110-131.
- de Vega, M. (1995). Backward updating of mental models during continuous reading of narrative. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 373-385.
- Wagener, M., & Wender, K. F. (1985). Spatial representations and inference processes in memory for text. In G. Rickheit and H. Strohnner, *Inferences in text processing*. Amsterdam: Elsevier Science Publishers.
- Wilson, S. G., Rinck, M., McNamara, T. P., Bower, G. A., & Morrow, D. G. (1993). Mental models and narrative comprehension: Some qualifications. *Journal of Memory and Language*, 32, 141-154.

Author Note

This research was supported by grants to Arthur Glenberg from the AFOSR (89-0367, and AASERT grant F49620-92-J-0310) and the University of Wisconsin Graduate School Research Committee. Experiments 5-7 served as part of Douglas Kramer's master's thesis. Requests for reprints should be directed to William Langston, Department of Psychology, Denison University, Granville, OH 43023, email: langston@cc.denison.edu.

Footnotes

1. All of the patterns of spatial layouts are available from the first author.
2. There was only one item recognition probe for each text. So, for a given text the probe could be either notice, not notice, spatial end, spatial middle, item three, or some item not in the text (a "no" probe).
3. The noticing effect reported in Table 2 is different from that reported here because the data in Table 2 combine notice and not-notice probes from items one and two whereas the data in the text separate items one and two.
4. Thanks to Morton Ann Gernsbacher and Karen Luh for suggesting many of the novel features of this experiment.

Table 1

Sample Texts From Experiments 1 and 2

4 Item Notice version from Experiment 1

Sam was setting up his fish aquarium decorations.

Sam put the castle down first.

Next, he put the plastic diver right of the castle.

Then Sam put the big rock under the plastic diver.

Then Sam put the seaweed left of the big rock.

{Probe: castle}

Next, he put the treasure chest left of the seaweed.

Finally, Sam put the sunken boat over the treasure chest.

When Sam put the fish in, it was scared by all of the stuff.

4 Item Far Notice version from Experiment 2

Sam was setting up his fish aquarium decorations.

Sam put the castle down first.

Next, he put the plastic diver way right of the castle.

Then Sam put the big rock under the plastic diver.

Then Sam put the seaweed way left of the big rock.

{Probe: castle}

Next, he put the treasure chest left of the seaweed.

Finally, Sam put the sunken boat over the treasure chest.

Table 2
Summary of the Experiments

Exp	DVA	Modality	Item/Dim ^b	Picture	Notice ^c Effect	t med ^d		t med ^e		tz ^f		tz ^g Loose
						Standard	Loose	Standard	Loose	Standard	Loose	
1	Recognition	visual	4/2	no	21.7	1.33 (41)	1.20 (41)	.72 (36)	.32 (41)			
			4/2	yes	17.3	-.29 (38)	.75 (39)	.11 (36)	.61 (39)			
			6/2	no	6.7	-.45 (41)	.28 (41)	-.32 (36)	.74 (41)			
			6/2	yes	28.0	.66 (38)	1.51 (39)	.65 (36)	.97 (39)			
2	Recognition	visual	4/2	no	-21.1	-.61 (27)	-.99 (29)	.37 (9)	.57 (29)			
			4/2	yes	10.8	<u>2.11 (27)</u>	.53 (29)	1.22 (17)	1.88 (27)			
			4/2	unlabeled	30.0	.04 (27)	1.13 (27)	.41 (17)	.89 (27)			
			4/2 (far)	no	-23.3	.80 (27)	-1.09 (29)	.16 (9)	-1.32 (29)			
			4/2 (far)	yes	-0.4	-1.03 (27)	-.02 (29)	.35 (17)	.27 (27)			
			4/2 (far)	unlabeled	-14.4	.13 (27)	-.68 (27)	.61 (17)	-1.17 (27)			
			6/2	no	54.4	<u>3.00 (27)</u>	1.35 (29)	.98 (9)	<u>2.43 (29)</u>			
			6/2	yes	19.9	-.65 (27)	.67 (29)	-.44 (17)	-.64 (27)			
			6/2	unlabeled	-45.7	-.76 (27)	-1.67 (27)	-1.18 (17)	-1.21 (27)			
			6/2 (far)	no	4.3	-.14 (27)	.15 (29)	.57 (9)	-.35 (29)			
3	Recognition	visual	4/2	no	7.3	.60 (17)	.23 (17)	.60 (17)	.13 (17)			
		visual	4/2	yes	36.1	.43 (17)	.92 (17)	-.08 (17)	-.09 (17)			
		auditory	4/2	no	-47.3	-.34 (17)	-.23 (17)	-.48 (17)	-.51 (17)			
		auditory	4/2	yes	45.4	.97 (17)	.67 (17)	-.19 (16)	-1.12 (16)			
4 ^h	Recognition	visual	3/1	no	-11.7		-.36 (22)		-1.10 (22)			
5	Reading	visual	4/2	no	1.8	1.09 (45)	.03 (50)	.72 (45)	1.26 (50)			
6 ⁱ	Reading	visual	4/2	no	36.9	.13 (36)	.65 (39)	.21 (36)	1.29 (39)			
7 ^j	Reading	visual	3/1 ^k	no	50.4	.72 (37)			1.18 (37)			
			3/1 ^l	no	34.9	.42 (37)			1.57 (37)			

Table 2 continued

- ^aDependent variable, either speeded recognition of a probe (Recognition) or time to read an about sentence (Reading).
- ^bNumber of items before the probe or about sentence, and the number of spatial dimensions required to arrange the items. "Far" indicates that the noticing effect was computed as the difference between the notice and far conditions (see description of Experiment 2 for details).
- ^cDifference between medians (in msec) for the not-notice and notice conditions in the loose analysis
- ^dValue of the t statistic (and sample size) for the medians in the standard analysis
- ^eValue of the t statistic (and sample size) for the medians in the loose analysis
- ^fValue of the t statistic (and sample size) for the z-scores in the standard analysis
- ^gValue of the t statistic (and sample size) for the z-scores in the loose analysis
- ^hData are collapsed across items 1 and 2; because there was no comprehension question, the standard analysis could not be performed.
- ⁱData are from the condition in which the target was item 1
- ^jBecause correct responding was forced, there is no distinction between standard and loose
- ^kData are from the condition in which the frame is on item 2
- ^lData are from the condition in which the frame is on item 3

Table 3

Sample text from Experiment 4

Text

Alan was arranging some modes of transportation according to how comfortable they seem to him.
Alan started with trains.

Then way in front of trains Alan placed planes because they seem more comfortable.

Then immediately in back of planes Alan placed cars because they seem less comfortable.

Probes

Notice: trains

End: planes

Item 3: cars

"No": buses

Verification sentences

trains seem more comfortable than planes (F)

trains seem more comfortable than cars (F)

planes seem more comfortable than cars (T)

planes seem more comfortable than trains (T)

cars seem more comfortable than trains (T)

cars seem more comfortable than planes (F)

Table 4

Sample notice text from Experiment 5

Kelly was arranging things on her night table.
She put the book down first.
Then she put the jewelry box to the right of the book.
Next, Kelly put the alarm clock in front of the jewelry box.
Then she put the photograph to the left of the alarm clock.
The book was a gift.
Finally, Kelly moved the book in front of the alarm clock.

Comprehension Question

How was the book obtained?

Table 5
Sample texts from Experiment 7

Introduction

Doug was setting up for a picnic.
First, he positioned the cake.

Notice Text, Frame of Reference on Item 3

Then he put the picnic basket to the right of the cake.
Next, he put the box of chicken directly to the left of the picnic basket.
So, the picnic basket is on the box of chicken's right.
The cake was strawberry.
The picnic basket broke.

Notice Text, Frame of Reference on Item 2

Then he put the picnic basket to the right of the cake.
Next, he put the box of chicken directly to the left of the picnic basket.
So, the box of chicken is on the picnic basket's left.
The cake was strawberry.
The picnic basket broke.

Not-Notice Text, Frame of Reference on Item 3

Then he put the picnic basket to the right of the cake.
Next, he put the box of chicken directly to the right of the picnic basket.
So, the picnic basket is on the box of chicken's left.
The cake was strawberry.
The picnic basket broke.

Not-Notice Text, Frame of Reference on Item 2

Then he put the picnic basket to the right of the cake.
Next, he put the box of chicken directly to the right of the picnic basket.
So, the box of chicken is on the picnic basket's right.
The cake was strawberry.
The picnic basket broke.

Test Questions

Is the cake anywhere to the right of the box of chicken?
Is the picnic basket anywhere to the left of the cake?
Was the cake chocolate?
Did the picnic basket break?

Robertson and Glenberg, full report:

Sketching does not facilitate learning from texts: A comparison of three study tasks
There is considerable evidence that visual illustrations can help readers understand and remember information from texts. Levin, Anglin and Carney (1987) reviewed 100 experiments on the benefits pictures have on learning from texts. They found an average improvement in comprehension of .71 standard deviations for readers of texts with text-relevant illustrations compared to readers of unillustrated texts. In a similar review of 46 experiments, the average effect size was .55 standard deviations (Levie & Lentz, 1982).

In this paper we ask a related question: Can readers enhance comprehension by drawing their own pictures based on the texts they are reading? We begin with a review of the theoretical reasons to believe the answer is yes, and follow this with a brief review of the empirical support. In fact however, after the results of several experiments, we conclude the answer is a qualified no.

While there can be no disagreement that illustrations can facilitate learning from texts, there are a number of conflicting theoretical approaches to explain the benefits of illustrations. The theoretical approach that we adopt and test in this paper is that of spatial mental models. A mental model (Johnson-Laird, 1983; Gentner & Stevens, 1983), or situational model (van Dijk & Kintsch, 1983), is a mental representation of the situation described in a text rather than a representation of the words themselves. During language processing, listeners and readers form mental models in order to understand the linguistic messages they receive. Any technique that would aid the formation of mental models should give benefits for comprehension and memory. One explanation for the beneficial effects of illustrations is that they promote the formation of mental models.

Given the general framework of mental models there are numerous ways to characterize their makeup. One particular characterization of mental models that we (Glenberg, Meyer, & Lindem, 1987; Glenberg & Langston, 1992; Glenberg, Kruley & Langston, 1994) and others (Denis & Cocude, 1989, 1992; Morrow, Greenspan & Bower, 1987; Morrow, Bower & Greenspan, 1989; Rinck & Bower, 1995) have advanced is that mental models used in language comprehension are spatial and picture-like in nature. For example, Denis and Cocude (1989, 1992) have shown that mental representations generated from hearing a description of a layout are similar to memory representations acquired from studying a map or directly scanning a map. Rinck and Bower (1995) found that spatial distance in readers' mental models formed by memorizing a spatial layout predicted reading times for sentences that referred to distant objects, relative to a character's described position. Glenberg, et al. (1987) similarly demonstrated an effect of spatial distance between a character and an object described in short texts. Glenberg and Langston (1992) proposed that spatial mental models could be used in situations which are only metaphorically spatial, such as the ordering of events in time. The studies cited above show that spatial relationships in a narrative can have effects on behavior independent of propositional structure or grammatical structure.

An important aspect of mental models is that they support inferences, allowing readers to learn more than just what is stated verbatim in a text. Making appropriate inferences is important in gaining a coherent representation of a situation. It allows comprehenders to learn and reason from what they are given. For example if you are told that Bob is taller than Jerry and that Kevin is taller than Bob, you can make the inference that Kevin is also taller than Jerry. This inference could be made using a mental model which arrays the people along a dimension of size.

This spatial characterization of mental models produces a clear prediction about drawing: Drawing a picture forces the reader to construct a spatial mental model, resulting in better comprehension and memory. This linking between a drawing and a drawer's mental representation occurs because the formation of a picture-like mental representation "employs inference processes like those that make information explicit in the course of drawing a diagram." (Larkin and Simon, 1987, p. 98) In essence, drawers can bootstrap their understanding of a text by creating a partial external representation of their current understanding of the text and using it to increase their understanding. In particular the prediction from spatial mental models is that both

illustrations and drawings should facilitate comprehension of information that is portrayed in them, and should facilitate inference-making through similar mechanisms.

There is some empirical evidence suggesting that drawing may be an effective strategy for learning from texts. In a study of adjunct study aids, Snowman and Cunningham (1975) demonstrated that drawing pictures was just as beneficial as answering adjunct questions. Their participants read a passage (from Anderson & Myrow, 1971) describing a fictitious African tribe, the Gruanda. Some participants wrote answers to questions, and others drew pictures in response to prompts. For example, some participants responded to the question "what do the Gruanda pay their taxes with" and other participants were instructed to "sketch a picture depicting what the Gruanda pay their taxes with." On a post-reading comprehension test, participants were asked multiple choice questions such as "what do the Gruanda use to pay their taxes." The comprehension scores were compared to scores of control-group participants who read the texts without accompanying tasks. The slight benefit of drawing was not statistically significant, and was found only for those test items that had been encountered during practice.

Dean and Kulhavy (1981) used the same materials with a different task. Some participants were specifically instructed to draw a map of the Gruanda territory and events described in the passage. Their performance was compared to participants who only read the text and did not get instructions for what information to emphasize. Draw participants performed significantly better than read-only participants both on measures of free recall and on comprehension questions.

Alesandrini (1981) had participants either read and summarize (which she termed paraphrase), read and draw, or read-only, while studying a science chapter describing the structure and functioning of batteries. After each new concept was introduced, draw and summarize participants were prompted to draw or summarize the new ideas. The read-only participants were told to read the text twice. On a comprehension test, both draw and summarize participants performed better than read-only participants, with draw participants doing better than summarize participants. Alesandrini also assigned additional study strategies (holistic or analytical) in combination with the drawing task. These strategies served to further enhance performance.

These three studies provide some evidence that guided drawing can facilitate learning from texts relative to read-only controls. There have been similar studies with similar results using children (e.g. Gobert & Clement, 1994; Gobert, 1995). However, the cause of the facilitation is not clear. The benefits attributed to drawing may have resulted because drawing facilitates comprehension. Alternatively, the benefits may have resulted from the strategic guidance that was given to the experimental groups that was not given to the control condition participants. That is, in the draw conditions, participants were told what to draw and that material was tested subsequently. Conversely, in the control conditions, to-be-tested information was not highlighted. Also the amount of time each participant was exposed to the text was not controlled in these studies, with read-only participants spending less time on task.

The hypothesis that drawing facilitates learning from text gains some support from empirical evidence that reader-generated mental images (i.e. pictures in the head) can facilitate learning. If readers are instructed to perform mental imagery while reading there are positive effects on memory for the text compared to reading without visualization, but the effect will be smaller than if an actual picture were provided. As summarized by Levie and Lentz (1982), the effects of mental imagery are found only for some types of readers in some situations. In particular it appears that the benefits of visual imagery on learning from texts are restricted to older, skilled readers. Levin et al. (1987) reported that the benefit of representational imagery was around one third of a standard deviation, "too small to be of any educational significance" (Levin et al., 1987, p. 71).

In contrast to the arguments predicting that drawing will facilitate learning from texts, it is possible to posit a theoretical explanation that supports the robust beneficial effects of illustrations, but does not predict that drawing would be a beneficial study activity. Consider the following claim: illustrations in texts aid readers because they contain information not readily acquired from texts alone. Studies that examine the effects of illustrations on text comprehension generally claim

that the illustrations used in the experiments overlap the texts in content so that the illustrations do not present new information to the readers that is not in the texts. But this does not necessarily imply that the benefits of illustrations occur because the mental representation is picture-like as the mental models approach would have us believe. Larkin and Simon (1987) demonstrate that diagrammatic descriptions (which include pictures) are computationally superior to verbal descriptions for communicating and supporting problem-solving, regardless of the nature of mental representation.¹ Diagrams promote faster searches, reduce the need for symbolic label matching, and support automatic inferences made with the perceptual system. So it is possible that illustrations promote comprehension solely on the basis of computational superiority and bear no resemblance to internal mental representations.

Given this claim about the cause of the benefits of illustrations, there is no reason to expect that drawing would help, because readers first must comprehend the information from the text in order to draw it. Similarly, if readers have not yet grasped important information from the text, then they will not draw it. By this reasoning, drawing would not facilitate learning beyond providing motivation for careful reading. Drawing may actually hinder comprehension by diverting attentional resources.

Our goals in this research are to see if there is a benefit of drawing, and if there is a benefit, to see if the benefit is for particular types of information. Spatial mental model theory predicts that drawing would facilitate memory for pictured information and inference-making. The spatial mental model theory advanced by Glenberg and Langston (1992) and Glenberg, Langston and Kruley (1994) predicts that the benefits should extend to situations that are literally and metaphorically spatial.

Our experiments test these hypotheses by assigning study tasks to readers and testing their comprehension. Our experiments differ from the studies mentioned previously (Alesandrini, 1981; Dean & Kulhavy, 1981; Snowman & Cunningham, 1975) by addressing the effects of drawing irrespective of other strategies: We provided similar strategic guidance to participants in each condition. In addition, each of the previous studies used a single text, so it is impossible to generalize the findings beyond that text. Our study uses multiple exemplars of two types of texts, spatial and non-spatial. We also manipulate the types of comprehension questions to determine if drawing benefits some types of information more than other. Finally, we employ an experimenter-paced task, rather than subject-paced, to ensure that time on task is at least nominally similar across conditions.

Experiment 1

We designed three activities similar to the ones used by Alesandrini (1981). Participants either read and summarized, read and drew, or read and thought. Since we are interested in the effects of the tasks, we kept all of the instructions as similar as possible across these experimental conditions.

Two types of texts were included in this study. The first type, spatial texts, described physical situations in which objects or characters were described as being arranged physically in space. These are types of texts most commonly used in studies of reading with pictures or reading with drawing, because they map easily onto a pictorial representation. Non-spatial texts described situations that were only metaphorically spatial. Use of non-spatial texts was motivated by research that shows that diagrams can aid comprehension of texts that describe non-spatial situations, for example the ordering of events in time (Glenberg & Langston, 1992; Small, Lovett, & Scher, 1993). It has been suggested that such non-spatial situations may be comprehended via metaphors to spatial situations (Lakoff, 1987), so we were interested to see if drawing while reading would have similar effects on these texts. Examples of spatial and non-spatial texts appear in Table 1.

insert Table 1 about here

Four types of comprehension questions arise from the factorial combination of two independent variables. The first variable is whether or not the queried information is likely to be pictured in a

drawing. Within any text, some items are easier to picture than others. For example, (1) From left to right she places the flour, the sugar, and the cocoa powder. is fairly easy to draw while (2) Anne learned to cook from her mother. is more difficult to draw. Consequently we thought that if readers attempted to draw the situation of Anne arranges the ingredients the way her mother taught her, placing the flour, sugar, and cocoa powder from left to right. they would likely draw something representing (1) and omit information like (2). There is empirical evidence that illustrations paired with texts facilitate learning pictured information better than non-pictured information (Levie & Lentz, 1982). Of interest is if this principle extends to reader-generated pictures. Obviously we were not able to directly manipulate whether our participants would draw the information or not; we could only manipulate our judgment of ease of picturing. We provide a manipulation check of this variable in the results below.

The second variable controlling the types of comprehension questions was the level of explicitness of the answers. The answers to some questions appeared verbatim in the texts, whereas other questions required some inference to be solved. Two issues motivate including this factor. First, it has been suggested, by Taylor and Tversky (1992), that verbatim questions are answered from a representation of the language of the text, whereas inference questions are answered from a mental model. So testing both types of information allows us to gauge which aspects of the texts are emphasized by a particular study task. Second, valid measures of reading comprehension should indicate a reader's ability to use the information in the text, not just parrot it back. For external validity, a study task should promote inference-making to be of educational value.

To summarize, the first experiment investigates the effects of three study tasks (summarize, draw, read-only) on memory for written texts. The content of the texts was predominantly spatial or non-spatial. Comprehension questions varied according to whether the answers were likely- or unlikely-to-be-pictured, and how explicitly the answers appeared in the texts. Importantly, we sampled several exemplars of the two types of text, and strategic guidance was comparable in all study conditions. Spatial mental model theory predicts a benefit of drawing for likely-to-be-pictured information and a benefit for questions requiring inferences.

Method

Participants.

Seventy-two undergraduate college students participated in the experiment for extra credit in an introductory psychology course. All participants were native English speakers. Participants were randomly assigned to one of the three study conditions.

Materials.

Eight experimental texts and one practice text were written. Four of the experimental texts described spatial environments (an arrangement of eight items in a kitchen, a description of a painting that has 5 areas, a description of a small town with four districts, and a description of two people buying things at a shopping mall). The other four texts described non-spatial environments (a marketing survey of toys, including a price/demand curve and rankings of different models of toys; a flood situation describing rising river levels that combine with tidal flow; livestock weight-gain that describes weight gain per week and the amount of feed consumed; and a candidate selection process that ranks candidates on two dimensions). The non-spatial texts all involved numerical or ordinal relationships. Texts ranged from 9 to 17 sentences in length and averaged 182 words in length. See Table 1 for examples of the texts and test questions. These texts are shorter than the studies investigating study tasks cited earlier and longer than the texts in the studies investigating mental models. The lengths were chosen to allow for multiple texts in the experiment (so not too long) and to allow for multiple types of questions for each text (so not too short).

There were eight questions for each text, two of each of the four types of questions. Questions varied in explicitness, being either verbatim or inference. Answers to verbatim questions were surface transformations of sentences in the text; inference questions required some reasoning from the material in the text. Answers to the questions were either likely to be pictured, or unlikely to be pictured. Items judged to be likely to be pictured were easy to draw.

Unlikely-to-be-pictured items were difficult to draw. In the results below we show that this manipulation was successful. Likely to be pictured questions were depicted in participants' drawings 82% of the time, whereas unlikely to be pictured questions were virtually never depicted.

Design and Procedure.

Participants performed one of three learning activities while reading the texts. They either summarized the text (summarize), drew a picture representing the text (draw), or just studied the text (read-only). All participants were told they would be reading texts and performing a learning activity for a given amount of time before answering comprehension questions from memory. Summarize participants were told to summarize the most important information from the text using complete sentences. Draw participants were instructed to draw pictures representing the most important information in the texts. They were told they could label parts of their pictures, but not to write phrases or sentences. Read-only participants were instructed to think about the most important information in the texts. They were not permitted to make any marks on their test packet. It was left to the participants in all groups to determine what was the most important information in each text.

Four minutes were allotted for participants to read each text and perform their learning activity. Three minutes were allotted to answer questions about the text from memory. Pilot testing showed that four minutes was around the average amount of time needed to read and summarize the texts. Drawing took less time than summarizing. Three minutes appeared to be more than enough time to attempt to answer the questions. Participants were warned when there were thirty seconds remaining in both reading and answering phases. Additionally, the practice text allowed participants to gauge what level of detail they could include in their drawings and summaries in the given amount of time.

The experiment was run in groups of four to ten participants. Each participant received a packet with the texts and comprehension questions stapled in counterbalanced order. The top sheet had the instructions for the task. Participants read the instructions, performed their learning activity while reading a sample text and answered a set of sample questions. During this practice period, the experimenter watched to ensure that participants were performing their tasks correctly. Participants were invited to ask questions before and after the practice period. After the task began, the experimenter monitored their performance to discourage looking back to the text during the answering phase. The entire task took about 75 minutes. There was a short break after the first four experimental texts had been read and tested.

Results

Answers to the comprehension questions were scored as correct (1 point), incorrect (0 points) or half-correct (1/2 point). Many questions had multiple parts; half credit was scored for partial answers. For example the answer to "What did Jeff buy at Sears?" is a hammer and a picture frame. Half credit was awarded to answers that listed just one of the items. The coder was blind to the participants' conditions during scoring. A summary of the data, reported as percentage correct, appears in Table 2.

insert Table 2 about here

Some of the results are consonant with previous findings and theory, while others are surprising. The surprising result is that participants in the draw condition did not outperform read-only participants, indeed they performed worse. The result that replicates previous research is that draw participants performed better than summarize participants. These results are summarized in Figure 1.

insert Figure 1 about here

Accuracy on the questions was analyzed using a mixed design ANOVA. The between-subjects factor was study task (summarize, draw, read-only). The within-subjects

factors were text type (spatial or non-spatial), explicitness (verbatim or inference), and pictured (likely- or unlikely-to-be-pictured). Two analyses were conducted, one on proportion correct and one on the arc sine of the square root of the proportion correct. The analyses were essentially identical, so the results from analyzing the untransformed scores are reported here.

Because there were four exemplars of each type of text (spatial and non-spatial) presented to each participant, it is possible to perform the analysis treating the replications of the texts as the random factor as well as treating subjects as a random factor. This allows us to check that the results generalize across the stimulus set as well as across subjects. For the items analysis, study task is manipulated within-items and text type is the between-items factor. The results of the analyses are reported here with the results by subject (subscripted with 1) followed by the results by item (subscripted with 2). Mean square errors (MSE), and measures of effect size are reported from analyses using subjects as a random factor.

Task and interactions with task are the comparisons of interest to our research question, so these results are presented first. There was a significant main effect of task, with summarize participants performing worst and read-only participants performing best [Means=.76, .80, .86, MSE=.071; $F(2,69)=5.79$, $F(2,12)=6.83$, $p<.01$]. Pairwise contrasts were computed using the error term relevant to the contrast, rather than using the pooled estimate (per Keppel, 1991, p. 384). Performance in the draw condition was significantly worse than the read-only group [MSE=.21, $F(1,46)=4.68$, $F(1,6)=7.80$, $p<.04$].² Performance in the summarize condition was also worse than in the read-only group [MSE=.30, $F(1,46)=10.91$, $F(1,6)=12.70$, $p<.02$]. The difference between the summarize and draw conditions was not statistically significant [MSE=.33, $F(1,46)=1.93$, $F(1,6)=1.94$, $p>.17$].

There was an interaction between task and explicitness [MSE=.0594 $F(2,69)=4.72$, $F(2,12)=5.09$, $p<.03$]. Inference questions were harder for all conditions, but the relative difficulty of inference questions versus verbatim questions was largest for the draw condition and least for the read-only condition (verbatim-inference: summarize=.08, draw=.13, read=.05). This is in the opposite direction predicted by spatial mental model theory. Pairwise contrasts demonstrated that the relative difficulty of inference questions was greater for draw participants than it was for read-only participants [MSE=.041, $F(1,46)=13.60$, $F(1,6)=8.34$, $p<.03$], but neither draw nor read-only participants differed significantly from summarize participants [F s all < 2.91, $p>.13$].

The interaction between task and pictured was marginally significant [MSE=.0783, $F(2,69)=2.47$, $p<.1$; $F(2,12)=5.26$, $p<.03$]. Questions that were about items that were likely to be pictured were harder in general, but not for draw participants (unlikely-likely: summarize=.08, draw=.02, read=.06). Pairwise contrasts revealed that the relative difficulty of pictured questions was greater for summarize participants than draw participants [MSE=.0795, $F(1,46)=4.72$, $p<.04$; $F(1,6)=11.52$, $p<.02$], but neither group significantly differed from the read-only participants. This was qualified by a significant three-way interaction between task, pictured, and text-types [$F(2,69)=3.76$, $F(2,12)=7.24$, $p<.03$]. The differences between likely- versus unlikely-to-be-pictured according to task depended on the type of the text. Draw and read-only participants did slightly better on likely to be pictured questions for spatial texts and considerably worse on non-spatial texts whereas summarize participants performed worse on likely to be pictured questions regardless of the text type. Summarize participants differed significantly from read-only participants [MSE=.081, $F(1,46)=5.50$, $F(1,6)=11.50$, $p<.03$] and from draw participants [MSE=.081, $F(1,46)=4.25$, $p<.05$; $F(1,6)=10.58$, $p<.02$], but draw and read-only participants did not differ in this respect (MSE=.050, F s < 1). These interactions with study task did not replicate in the second experiment, so it is not clear how confident we can be that they are real. However, across the experiments, participants in the draw condition appear to have a relative advantage for likely-to-be-pictured questions.

To check the validity of our distinction between likely- and unlikely-to-be-pictured questions we inspected the drawings produced by the first eight participants. Materials were coded for each question for whether the answer was apparent in the drawing, with the coder being blind to the participants' actual answers to the questions. The answers to likely-to-be-pictured

questions were identified in the drawings 81.6% of the time. We collated this data with the accuracy data to determine the strength of the relationship between including an item in a picture and being able to answer a question about it later. Not surprisingly, participants scored higher on questions for which they had included the relevant information in their drawings [Mean not-pictured=.57, s.d.=.21; Mean pictured=.89, s.d.=.08, $t(7)=3.51$, $p<.01$].

It is possible that participants were spending part of the reading phase trying to guess what questions would later be asked and answering them ahead of time. If this were the case, this could explain why read-only participants outperformed the participants with overt tasks, because the absence of demands for overt behavior would afford readers more time to plan ahead. If participants were trying to guess what questions would be asked during the study phase, they should progressively improve throughout the task as they better learn to anticipate the questions. To determine if there were detectable practice effects, another mixed ANOVA was performed using order of presentation of the eight texts as a factor, crossed with task, pictured, and explicitness (collapsed across text-type). There was no main effect of order [$F(7,483)=1.28$, $p>.25$], no interaction with task [$F(14,483)=1.21$, $p>.25$], and no higher-order interactions with task and other factors ($F_s<1$). We conclude that any practice effects were slight.

For completeness we report a number of results that do not involve the factor of study task. There was a main effect of explicitness of question, with verbatim questions being easier than inference questions [means .85, .77, $MSE=.070$, $F(1,69)=66.3$, $F(1,6)=17.36$, $p<.01$]. There was a main effect of pictured (only significant by subjects), with unlikely-to-be-pictured being easier than likely-to-be-pictured [means .83, .78, $MSE=.048$, $F(1,69)=20.4$, $p<.01$; $F(1,6)=2.40$, $p>.1$]. Pictured and explicitness interacted [verbatim-likely=.86, inference-likely=.70, verbatim-unlikely=.84, inference-unlikely=.83, $MSE=.013$, $F(1,69)=69.6$, $F(1,6)=12.44$, $p<.02$], so that the effect of explicitness only held for questions that were likely-to-be-pictured. Explicitness also interacted with text-type, so that whereas inference questions were always harder, this effect was enhanced on non-spatial texts [spatial/verbatim=.82, spatial/inference=.80, non-spatial/verbatim=.88, non-spatial/inference=.73, $MSE=.016$, $F(1,69)=40.1$, $p<.01$; $F(1,6)=11.42$, $p<.02$]. Finally, there was a significant three-way interaction between the factors of pictured, explicitness, and text type [spatial/non-spatial for following: verbatim-likely=.83/.90, verbatim-unlikely=.81/.86, inference-likely=.79/.61, inference-unlikely=.82/.84, $MSE=.012$, $F(1,69)=34.6$, $p<.01$; $F(1,6)=5.95$, $p<.051$].

Discussion

Two findings are notable. First, as others have found, drawing is a more effective study task than verbal summarizing, though this effect was not statistically significant in this experiment. Second, unlike others, we found that drawing does not facilitate performance compared to a read-only condition.

An explanation for the second finding is that readers must already have a coherent understanding of the text in order to draw useful pictures. But, if they already understand the text, then drawing is superfluous, and if they do not understand the text, then the drawings will not be useful. Furthermore, if the act of drawing diverts attention from comprehending, drawing may hinder performance.

Participants in the draw condition performed relatively better on questions coded as likely-to-be-pictured for spatial texts. This suggests that these questions tapped information that the participants emphasized while drawing their pictures. However, it should be noted that draw participants averaged three percent worse on these items than read-only participants. So drawing did not provide an advantage on these items, it just did not hinder them. In addition, the fact that draw participants were more accurate on questions tapping information they actually illustrated suggests that drawing may help for certain types of information depending on what participants include in their pictures.

For the most part, predictions derived from spatial mental model theory were not supported. The significant main effect of study task was in the opposite direction of the prediction. The interaction with explicitness, was also in the opposite direction of the prediction.

The only slight support for prediction based on spatial mental model theory was that drawing enhanced memory for pictured information.

The other notable finding for this experiment is the superior performance in the draw condition compared to the summarize condition. This is the same result that Alesandrini (1981) reported, and of similar effect size though it was not statistically significant here. At this point it is helpful to express the results of this experiment in the form of a standard effect size to allow us to compare results, as effect sizes are more useful in judging replication successes than statistical significance level (Rosenthal, 1990). Following Cohen (1988, p. 20), effect size, d , is computed as the difference between the two treatment sample means divided by the pooled estimate of the within-group standard deviation. The advantage of draw participants over summarize participants is medium to small in effect, $d=.23$. This is very similar to the effect size observed between draw and summarize conditions in Alesandrini's (1981) findings, $d=.28$. The effect size of the advantage of the read-only over the draw condition is also medium to small in effect, $d=.29$.

Experiment 2

The second experiment was designed to replicate and extend the results from the first experiment by including both immediate and delayed tests. All of the previously mentioned studies that manipulated drawing study tasks used only immediate tests (Alesandrini, 1981; Dean & Kulhavy, 1981; Snowman & Cunningham, 1975). However, long term retention is important for any practical applications concerning effects of drawing on memory. Also a delayed test with distracting material intervening should prevent participants from retaining a surface-form memory of the text through rote rehearsal, forcing them to answer from a long term representation. In this experiment the delay (filled with material from other texts) between reading and testing averaged twelve minutes. In a later experiment we extend this delay even further.

Method

Participants.

Eighty-one college students participated in the experiment for extra credit in an introductory psychology class. All participants were native English speakers. Each participant was randomly assigned to one of the three study tasks.

Materials.

This experiment used the same eight texts that were used in Experiment 1. The test questions were divided into two sets of four questions for each text so that one of each type of question was tested at each retention interval. The assignment of question sets to retention interval was counterbalanced. There was only one set of practice questions immediately after the practice text.

Procedure.

As in the first experiment, participants read the texts and performed their assigned learning activities. After reading each text (for four minutes), participants immediately answered comprehension questions for that text (for 1.5 minutes). After reading and answering questions for four texts, participants answered another set of comprehension questions for those four texts. Participants were given 1.5 minutes to answer the delayed questions for each text. Thus the delayed questions occurred an average of twelve minutes after the end of the reading phase, with considerable activity intervening, making deliberate rehearsal unlikely. As in the first experiment, after the first four texts had been read and tested, there was a short break. Reading, immediate testing, and delayed testing proceeded in the same manner for the final four texts.

The instructions were modified slightly from the first experiment to stress that the intended purpose of the learning activities was to facilitate memory for the texts. Participants were asked to make use of their study activity in whatever way they thought would be most beneficial, but to restrict themselves to the assigned activity.³ The instructions also stressed that each text would be tested twice. The experimenter monitored participants' performance on the tasks during the experiment to discourage looking ahead or behind in the test packets.

Results

Answers to the comprehension questions were scored as in Experiment 1. The coder was blind to the participants' conditions during scoring. A summary of the data appears in Table 3 and

Figure 2.

 insert Table 3 and Figure 2 about here

As in Experiment 1, read-only participants performed better than draw participants who performed better than summarize participants. The differences between the task conditions is primarily in the immediate test. On the delayed test, performance in the draw and read-only conditions was similar. Performance in the summarize condition was slightly worse.

Accuracy on the questions for each text was analyzed as a mixed design ANOVA. The between subjects factor was study task (summarize, draw, read). The within subjects factors were text type (spatial or non-spatial), retention interval (immediate or delayed), explicitness (verbatim or inference), and pictured (likely- or unlikely-to-be-pictured). Task and interactions with task are the comparisons of greatest interest. There was a marginally significant main effect of task, with performance in the summarize condition (.78) worse than the draw condition (.82) which was worse than the read-only (.84) condition [$MSE=.124$; $F(2,78)=3.06$, $p<.053$, $F(2,12)=6.70$, $p<.02$].

Predictably, there was a main effect for retention interval with scores higher for the immediate test than the delayed test [means .84, .79, $MSE=.037$, $F(1,78)=16.06$, $F(1,6)=24.32$, $p<.01$]. Retention interval did not significantly interact with task [$F(2,78)=1.80$, $F(2,6)=2.03$, $p>.17$] or combinations of task and other factors. For the rest of this analysis, we break the data down into simple effects according to retention interval.⁴

For the immediate test, there was a significant main effect of task, with performance in the summarize condition (.81) worse than the draw condition (.83) which was worse than the read-only condition (.87) [$MSE=.28$; $F(2,78)=3.23$, $F(2,12)=5.45$, $p<.05$]. This effect replicates the results from Experiment 1, however the differences among the means are about half the size. The relatively small effect may be caused by using only half the number of questions used in Experiment 1. Pairwise contrasts reveal that performance in the read-only condition was marginally better than the draw condition [$MSE=.19$, $F(1,52)=3.47$, $p<.07$; $F(1,6)=9.85$, $p<.02$]. Performance in the read-only condition was also higher than the summarize condition [$MSE=.34$, $F(1,52)=5.38$, $F(1,6)=9.39$, $p<.03$]. However the summarize and draw groups did not differ ($F_s < 1.14$, $p>.3$).

There was an interaction (significant only by subjects) between task and text-type. Both draw and read-only participants performed better on spatial texts than non-spatial texts, whereas summarize participants performed similarly on both types of texts [$MSE=.13$, $F(2,78)=3.93$, $p<.03$, $F(2,12)=3.04$, $p<.09$]. The only significant pairwise interaction contrast is between the summarize and draw groups [$MSE=.153$, $F(1,52)=6.69$, $p<.02$; $F(1,6)=4.15$, $p<.09$]. There was also a significant three-way interaction between task, explicitness and pictured [$MSE=.10$, $F(2,78)=3.81$, $F(2,12)=3.64$, $p<.03$]. In general, inference questions were harder than verbatim questions, and items that were likely-to-be-pictured were harder than those that were unlikely-to-be-pictured. However, for draw participants the likely-to-be-pictured items were actually easier than the unlikely-to-be-pictured, but only for the verbatim items; whereas read-only participants performed the same on these questions, and summarize participants performed worse on likely-to-be-pictured, regardless of explicitness. Draw and read-only participants differed in how they performed according to explicitness and pictured [$MSE=.10$, $F(1,52)=5.56$, $F(1,6)=10.81$, $p<.03$], draw and summarize differed [$MSE=.113$, $F(1,52)=5.26$, $F(1,6)=12.09$, $p<.03$], but summarize and read-only did not differ significantly ($F_s < 1$). Neither of these interactions was significant in the other experiments.

The manipulation of study task had a very small effect on the delayed test. The main effect of task was significant by items, but not by subjects. [$MSE=.36$, $F(2,78)=2.41$, $p<.1$, $F(2,12)=4.58$, $p<.04$]. Performance in the draw condition (.81) was similar to the read-only condition (.80) which was better than the summarize condition (.75), though not by much relative to the variability. A pairwise contrast on the extreme groups, draw and summarize, was significant only by items [$MSE=.43$, $F(1,52)=3.52$, $p<.07$; $F(1,6)=7.13$, $p<.05$]. There were

no interactions between study task and any of the other factors on the delayed test [F 1s all <1.11 , F 2s all <1].

Another mixed ANOVA was performed using order of presentation as a factor, crossed with task, explicitness, retention interval and pictured (collapsed across text-type). There was no main effect [$F(7,546)<1$], no interaction with task [$F(14,546)<1$], and no interactions with other factors (F s all <1.57 , p s all $>.14$). Thus there was little evidence for changes across the course of the experiment.

There are a number of results that do not involve the factor of study task that are reported here for completeness. Only effects that exceeded a $p=.05$ cutoff for both by-subjects and by-items analyses are reported. The factors of text type and explicitness interacted, with the factor of explicitness only affecting non-spatial texts [spatial/verbatim=83, spatial/inference=84, non-spatial/verbatim=85, non-spatial/inference=73, $F(1,78)=40.84$, $F(1,6)=7.13$, $p<.04$]. Explicitness also interacted with pictured, such that inference questions that were likely-to-be-pictured were much harder to answer [verbatim-likely=84, verbatim-unlikely=84, inference-likely=73, inference-unlikely=84; $F(1,78)=42.08$, $F(1,6)=6.14$, $p<.05$]. These two-way interactions were involved in a three way interaction between text type, explicitness and pictured. This interaction reveals that inference questions that were likely-to-be-pictured only differed from other questions on non-spatial texts [spatial/non-spatial mean percentages: verbatim-likely=83/86, verbatim-unlikely=83/84, inference-likely=85/61, inference-unlikely=84/85; $F(1,78)=61.95$, $F(1,6)=7.8$, $p<.04$]. The main effect of retention interval interacted with the factors of explicitness and pictured. This interaction reveals that there was more forgetting of unlikely-to-be-pictured verbatim information and of likely-to-be-pictured inference information [immediate/delayed mean percentages: verbatim-likely=85/83, verbatim-unlikely=87/81, inference-likely=76/69, inference-unlikely=86/83; $F(1,78)=4.67$, $F(1,6)=6.07$, $p<.05$]. As in Experiment 1, these factors that do not involve the manipulation of study task are not of theoretical importance to this research question, and will not be discussed further.

Discussion

This experiment successfully replicated the effects observed in the first experiment and extended them to delayed tests. The ordering of the conditions for the immediate test point was the same in both experiments. The order of the means was summarize (lowest), draw, and read-only. On the delayed test summarize participants performed worse than read-only participants who were roughly the same as draw participants. Importantly, in neither experiment was there evidence for benefits from drawing as a study task. Thus the results for the first two experiments go against the predictions from spatial mental model theory. The effects, where significant, are in the wrong direction from the predictions.

The higher-order interactions with task that were significant in Experiment 1 were not significant for Experiment 2. This is not too surprising as the effects were fairly small in Experiment 1, and there were fewer observations per cell in Experiment 2. However, it does appear that draw participants have a slight relative advantage for likely-to-be-pictured items, though this advantage generally is not adequate to boost their performance above read-only participants. This is evident in Experiment 2 in the relative advantage draw participants had for verbatim likely-to-be-pictured items.

Experiment 3

The results in the first two experiments that draw participants did not outperform read-only participants is of interest because experimenter-provided illustrations typically do enhance learning (Levie & Lentz, 1982; Levin, Anglin & Carney, 1987). Perhaps, however, our texts are peculiar so that neither participants' drawings nor experimenter-provided illustrations would enhance learning. That the finding was replicated and generalized across eight texts alleviates this concern to some extent, but not fully. In Experiment 3 we compared a read-only condition to one in which participants read texts with accompanying illustrations. If we find the usual enhanced performance with illustrations, then we can be confident that our materials are not peculiar.

Method

Participants.

Sixty-six college students participated in the experiment for extra credit in an introductory psychology class. All participants were native English speakers. Each participant was randomly assigned to one of the experimental conditions. The data from two participants were discarded because they failed to follow directions.

Materials.

This experiment used the same eight texts that were used in Experiment 2. An illustration was generated for each text, using computer-aided drawing and graphing software. The illustrations were based on the information in the texts so that the texts and illustrations overlapped in content. Additionally the illustrations were made so that questions that had been coded as likely-to-be-pictured were actually pictured and the answers to unlikely-to-be-pictured questions were not in the illustrations.

Procedures.

Participants were instructed to read the texts and think about the most important information. Half of the participants received unillustrated texts, the other half received texts with pictures printed below them. The timing and distribution of activities (reading, immediate test, and delayed test) were identical to Experiment 2.

Results

Answers to the comprehension questions were scored as in the first experiments. Additionally, the data from six participants were independently scored by another coder. The correlation between the two coders was .93.

Accuracy was analyzed as a mixed design ANOVA. The between-subjects factor was illustrated or not-illustrated. The within subjects factors were text type (spatial or non-spatial), retention interval (immediate or delayed), explicitness (verbatim or inference), and pictured (whether the answers to questions were in the illustrations or not). A summary of the data appear in table 4.

insert Table 4 about here

There was no main effect of illustrations, though performance was marginally better in the illustrated condition [mean accuracy for illustrated=.80, unillustrated=.78, $MSE=.167$; $F(1,62)<1$, $F(1,6)=2.18$, $p>.15$]. There were two interactions with illustrations. Readers of illustrated and unillustrated texts performed similarly on unpictured questions (.78 and .79), but readers of illustrated texts performed better on pictured questions (.80 and .75). This interaction was statistically significant only by-subjects [$MSE=.030$, $F(1,62)=6.80$, $p<.02$; $F(1,6)=2.75$, $p<.15$]. Readers of illustrated and unillustrated texts performed similarly on verbatim questions (.82 and .84), but readers of illustrated texts performed better on inference questions (.75 and .69), [$MSE=.033$; $F(1,62)=12.06$, $p<.01$; $F(1,6)=6.70$, $p<.05$].

Even though there was no interaction between retention interval and illustrations ($F_s<1$), it appears that the manipulation of illustrations had stronger effects for the longer retention interval. At both retention intervals participants with illustrated texts performed better on inference questions. But the interaction of illustrations with pictured questions was only significant by simple effects on the delayed test [For the immediate test both $F_s<1$; for the delayed test, $F(1,62)=8.03$, $p<.01$, $F(1,6)=7.01$, $p<.04$].

Discussion

This experiment replicates the well-known finding that illustrations facilitate text comprehension. Readers of illustrated texts are better able to answer comprehension questions than readers of unillustrated texts. In particular, readers of illustrated texts are better able to answer questions that require an inference from the text than are readers of unillustrated texts. Thus, we conclude that robust beneficial effects of illustrations extend to the stimuli in these experiments.

Experiment 4

Experiment 4 was designed to investigate a possible trend that appeared in the data of

Experiment 2. The data from the delayed test showed a possible switch in the order of our conditions, in that the sample mean for the draw condition was slightly but not significantly higher than the mean for the read-only condition. Perhaps with an even longer retention interval the draw condition will out perform the read-only condition.

We designed the next experiment to explore this trend by setting an even longer retention interval. Participants read all eight texts, followed by a distracter task and short break before being tested for comprehension. We also included a read-with-pictures condition (like that used in Experiment 3) to compare to the draw condition.

Method

Participants.

Eighty college students participated in the experiment for extra credit in an introductory psychology class. All participants were native English speakers except for one who performed comparably to the other participants. Each participant was randomly assigned to one of the experimental conditions.

Materials.

This experiment used the same texts and questions that were used in Experiment 1. For the read-illustrated condition, the pictures used in Experiment 3 were presented along with the texts. Texts were arranged into packets in counterbalanced order. Following the texts was a filler task unrelated to the materials or topic of this project. Lastly, sheets with questions for the texts were included in the packets.

Design and Procedure.

Participants were assigned to one of four experimental conditions: read and summarize, read and draw, read-only, and read-illustrated. Participants read the texts and performed their assigned learning activities for four minutes for each text. After all eight texts had been read and studied, participants performed a distracter task and had a short break. This filler period lasted from about six to nine minutes. Following the filler period, participants answered the questions about the texts from memory. In order to keep the time between texts and tests constant, the amount of time allotted to answer the questions was four minutes, the same duration as the study period. Consequently the delay between the presentation of a text and the answering phase for that text was approximately 40 minutes.

A practice text that was immediately followed by sample questions preceded the experimental texts. During the practice the experimenter watched to ensure participants understood their tasks. During the experiment the experimenter monitored performance to discourage looking ahead or behind in the test packets.

Results

Read-only participants performed better than participants with overt tasks. The potential trend observed in Experiment 2 did not materialize. A summary of the data appears in Table 5 and Figure 3.

insert Table 4 and Figure 3 about here

Answers to the comprehension questions were scored as in Experiment 1. The coder was blind to the participants' conditions during scoring. Accuracy on the questions was analyzed with mixed design ANOVA. The between subjects factor was study task (summarize, draw, read, read-illustrated). The within subjects factors were text type (spatial or non-spatial), explicitness (verbatim or inference), and pictured (likely- or unlikely-to-be-pictured).

Accuracy in the summarize and draw groups was very similar (both .70) and lower than in the read-only and read-illustrated groups (.77, .78). The main effect of task assignment was significant by items and marginally significant by subjects [$MSE=.125$, $F(3,74)=2.66$, $p<.054$, $F(3,18)=14.37$, $p<.001$]. A contrast grouping the overt tasks (summarize and draw) and the read-only tasks (read and read-illustrated) revealed these groups differ, [$F(1,74)=7.25$, $F(1,6)=57.86$, $p<.01$].

The only interaction with task assignment that was statistically significant (and only

significant by items) was the interaction with pictured [$MSE=.0316$, $F(3,76)=2.26$, $p<.09$; $F(3,18)=3.77$, $p<.03$]. Each condition except for the read-illustrated condition performed worse on questions that were likely-to-be-pictured than unlikely-to-be-pictured whereas the read-illustrated participants performed better on questions that were pictured than on questions that were not illustrated. Contrasts revealed that participants in the illustrated condition performed better than other participants on the illustrated questions [$F(1,76)=5.68$, $p<.03$; $F(1,18)=16.01$, $p<.001$], but not significantly better on unillustrated questions ($F_s < 1$), replicating the major result from Experiment 3. Conversely, draw participants did not exhibit this relative advantage for pictured items. An ANOVA excluding the read-illustrated group revealed no interactions of task with pictured ($F_s < 1$).

There are two results that do not involve the factor of study task that are reported here for completeness. The factors of explicitness and pictured interacted [$F(1,72)=63.6$, $F(1,6)=16.46$, $p<.01$]. Participants actually performed better on the unlikely-to-be-pictured inference questions than on the unlikely-to-be-pictured verbatim questions [means = .72 and .77], with the opposite pattern for the likely-to-be-pictured questions [.73 and .68]. This was mediated by a three way interaction between explicitness, pictured and text-type [$F(1,72)=25.29$, $F(1,6)=6.15$, $p<.05$; spatial/non-spatial for following: verbatim-likely=.77/.81, verbatim-unlikely=.71/.73, inference-likely=.80/.55, inference-unlikely=.80/.74, $MSE=.015$].

Discussion

This experiment successfully lays to rest any concerns that the pattern of results observed earlier would change as a function of retention interval. Even with a long retention interval, read-only participants outperformed the draw and summarize participants.

General Discussion

These experiments provide a striking result. Draw participants fail to outperform read-only participants. In fact, they performed worse. This result generalized over many texts and three experiments. This finding is in opposition to the predictions made from spatial mental model theory, and it challenges the conclusions of previous research. Thus, it requires an explanation.

One simple explanation is that the benefits of drawing or summarizing require long texts and the benefits are reversed for short texts. Our texts averaged 182 words in length, compared to 2190 words for the Gruanda text (Anderson & Myrow, 1971). This explanation, while perhaps lacking in parsimony, is tractable and the critical experiments should be straight-forward to perform. ⁵

Another explanation is that drawing per se does not aid readers, because they must already have a coherent understanding of the text in order to draw useful pictures. By this account, drawing pictures is redundant at best, and potentially distracting because it consumes resources that could have been used to comprehend the text. Why then did drawing appear to help comprehension in the studies reported previously (Alesandrini, 1981; Dean & Kulhavy, 1981; Snowman & Cunningham, 1975)? In those experiments, draw participants were given guidance as to what and how to draw, while read-only controls were just told to read the text. So, those results may reflect strategic information given to draw participants but not read-only controls.

The experiments reported here support the idea that illustrations in texts are beneficial because pictures present information that readers would not fully comprehend otherwise, despite the apparent overlap of information. Illustrations are beneficial because they provide computational advantages for scanning speed and recognition, and promote perceptual inferences (Larkin & Simon, 1987), not because they directly map onto a picture-like mental model. If spatial mental models were routinely used in comprehension, the act of drawing should have caused our draw participants to form more elaborate mental models and outperform the read-only participants. However, constructing detailed spatial mental representations of a situations described in texts is a difficult process that may not be accomplished without illustrations or special training. This argument is supported by experimental evidence that readers often do not make spatial inferences or do not make inferences unless their task explicitly requires it (Langston, Kramer & Glenberg, 1995; Taylor & Tversky, 1992).

After considering the results of these experiments, we reject the proposition that people comprehend texts by forming picture-like spatial mental models. We do not deny the considerable evidence that something like mental models exist, just that they are not of a picture-like nature. Instead mental models may reflect functional relationships. For example, Lakoff (1987) characterizes cognitive models as being composed of image-schemas, such as container, source-path-goal, link, and others. These image schemas provide structural, functional, and logical information. For example, the container schema is structured by a boundary into an interior and an exterior. Functionally, anything in the container is under the control of the container. Logically, something can be in or out of the container, but not both.

Although we found no direct evidence that drawing is an effective study task, there may be circumstances in which drawing does help. Our own tendency to sketch complicated interactions when reading reports of experiments suggests that drawing helps sometimes. We propose the following hypothesis: Readers will benefit from drawing provided, a) the text is complicated, and, b) the readers have domain-specific strategies relevant to producing beneficial pictures. The text must be complicated enough so that it is necessary to have external support to form a coherent representation. In these instances diagrams may help to reduce the computational load of sentential representations (Larkin & Simon, 1987). Furthermore, given a difficult text, readers must have appropriate drawing strategies for that text. For example, a complex interaction is unlikely to lead a novice reader to draw an appropriate diagram. This reasoning suggests that any facilitation by drawing will be domain specific, with minimal transfer to new situations.

Our hypothesis that drawing facilitates learning only when drawers have appropriate drawing strategies is supported by some evidence in the problem-solving literature. Kindfield (1993-4) demonstrated that experts draw more pictures, have higher quality pictures, and use them more effectively than novices. Perhaps one component of expertise is learning effective strategies for producing and using drawings.

In conclusion, our results demonstrate that drawing is generally not an effective study task: Reader-generated pictures do not provide the same benefits as experimenter-provided pictures. Based on these results, we reject the spatial mental models explanation for the benefits of pictures because it implies that drawing and pictures should produce similar benefits. Finally, we suggest that drawing may produce positive benefits if readers have strategies to help them produce relevant drawings for difficult texts, but this suggestion remains to be tested.

References

- Alesandrini, K. L. (1981). Pictorial-verbal and analytic-holistic learning strategies in science learning. Journal of Educational Psychology, 73, 358-368.
- Anderson, R. C., & Myrow, D. (1971). Retroactive facilitation in prose. Journal of Educational Psychology, 62, 354-378.
- Cohen, J. (1988). Statistical power analysis in the behavioral sciences (2nd ed.). Hillsdale, NJ: Erlbaum.
- Dean, R. S. & Kulhavy, R. W. (1981). Influence of spatial organization in prose learning. Journal of Educational Psychology, 73, 57-64.
- Denis, M. & Cocude, M. (1989). Scanning visual images generated from verbal descriptions. European Journal of Cognitive Psychology, 1, 293-307.
- Denis, M. & Cocude, M. (1992). Structural properties of visual images constructed from poorly or well-structured verbal descriptions. Memory and Cognition, 20, 497-506.
- Gentner, D. & Stevens, A. L. (Eds.). (1983). Mental models. Hillsdale, NJ: Erlbaum.
- Glenberg, A. M. & Langston, W. E. (1992). Comprehension of illustrated text: Pictures help to build mental models. Journal of Memory and Language, 31, 129-151.
- Glenberg, A. M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. Journal of Memory and Language, 26, 69-83.
- Gobert, J. (January, 1995). Diagramming versus summarizing during reading: The effects on the propositional and mental model representations. Presentation at the Sixth Annual Winter Text Conference. Jackson Hole, WY.
- Gobert, J. & Clement, J. (April, 1994). Promoting causal model construction in science through student-generated diagrams. Presentation at the Annual Meeting of the American Educational Research Association. New Orleans, LA.
- Johnson-Laird, P. N. (1983). Mental models. Cambridge, MA: Harvard University Press.
- Keppel, G. (1991). Design and analysis: A researcher's handbook (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Kindfield, A. C. H. (1993-4). Biology diagrams: Tools to think with. The Journal of the Learning Sciences, 3, 1-36.
- Lakoff, G. (1987). Women fire and dangerous things. Chicago: University of Chicago Press.
- Langston, W. E., Kramer, D. C., & Glenberg, A. M. (1995). Mental models are not (very) spatial. Manuscript in preparation. University of Wisconsin-Madison.
- Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. Cognitive Science, 11, 65-100.
- Levie, W. H., & Lentz, R. (1982). Effects of text illustrations: A review of research. Educational Communication and Technology Journal, 30, 195-232.
- Levin, J. R., Anglin, G. J., & Carney, R. N. (1987). On empirically validating functions of pictures in prose. In D. M. Willows & H. A. Houghton (Eds.), The Psychology of Illustration, volume 1, Basic Research. (pp. 51-85). New York: Springer-Verlag.
- Morrow, D. G., Greenspan, S. L., & Bower G. H. (1987). Accessibility in situation models in narrative comprehension. Journal of Memory and Language, 26, 165-187.
- Morrow, D. G., Bower G. H., & Greenspan, S. L. (1989). Updating situation models during narrative comprehension. Journal of Memory and Language, 28, 292-312.
- Rinck, M. & Bower, G. H. (1995). Anaphora resolution and the focus of attention in situation models. Journal of Memory and Language, 24, 110-131.
- Rosenthal, R. (1990). Replication in behavioral research. In Neuliep, J. W. (Ed.) Handbook of replication research in the behavioral and social sciences. [Special Issue.] Journal of Social Behavior and Personality, 5(4), 1-30
- Small, M. Y., Lovett, S. B., & Scher, M. S. (1993). Pictures facilitate children's recall of unillustrated expository prose. Journal of Educational Psychology, 85, 520-528.
- Snowman, J & Cunningham, D. J. (1975). A comparison of pictorial and written adjunct aids in learning from text. Journal of Educational Psychology, 67, 307-311.
- Taylor, H. A., & Tversky, B. (1992). Spatial mental models derived from survey and route

descriptions. Journal of Memory and Language, 31, 261-292.

van Dijk, T. A., & Kintsch, W. (1983). Strategies of discourse comprehension. New York: Academic Press.

Footnotes

1 Larkin and Simon speculate that picture-like mental representations could be operated on in a manner similar to diagrammatic external representations, and suggest that building a picture-like representation would use the same processes as constructing an actual picture. However, their analysis is of the efficiency of diagrams as external representations.

2 Because there are three possible pairwise contrasts and only two degrees of freedom for the factor of task, it may be desirable to adjust for family-wise type 1 error rate. Using a Bonferroni method, the comparison between the draw and read-only condition is significant at the $p=.06$ level, rather than the customary $p=.05$. This effect replicated in Experiments 2 and 4, so we assert it is valid.

3 This slight modification was made so that participants in the assigned task conditions would not feel burdened to produce good summaries or drawings if they did not believe it would help them remember the texts better. We thank Nancy Denney and Joel Levin for this modification which was recommended in an early review of the project.

4 We analyze by simple effects even though the interaction was not significant for several reasons. First, it allows us to judge the partial replication of Experiment 1 by considering the immediate test data separate from the delayed test data. Second, there appear to be different patterns at the two retention intervals (both by a cursory examination of the ordering of the conditions, and by formal analysis as follows), so the nonsignificant interaction may be caused by inadequate power. Finally, simple effects analysis is equivalent (by sums of squares) to the full factorial analysis with the added benefit in this case of reducing violations of homogeneity of variance.

5 We thank Joel Levin for bringing this to our attention.

Table 1: Sample texts followed by questions

Example of a spatial text

Anne is going to bake a cake for her roommate. Anne learned to be very organized in the kitchen from her mother. She reads through the recipe first and collects all of the ingredients. On the left side of her counter going from left to right she places the flour, the sugar, and the cocoa powder. They are all in a row behind her mixing bowl. Right behind the dry ingredients, Anne assembles all of the wet ingredients. She lines up, left to right, the oil, vanilla, and enough egg substitute to equal 3 eggs. She puts her recipe book on the right side of her counter to allow quick reference. She double checks her ingredients and is ready to start mixing.

Explicitness & Pictured:

Verbatim, likely: What ingredient is to the right of the flour?

Verbatim, unlikely: For whom is Anne baking the cake?

Inference, likely: What ingredient is in front of the egg substitute?

Inference, unlikely: What flavor cake is Anne making?

Example of a non-spatial text

A fisherman, Tom, docks his boat near the mouth of a river. He is close enough to the ocean that he is significantly affected by the rise and fall of the tides. The tides rise and fall every 12 hours, taking 6 hours to fall and 6 hours to rise. The tides currently are in a pattern such that high tide occurs at noon and at midnight. At this time of year, the tides range six feet of difference between high and low tide, from minus 2 feet to plus 4 feet, relative to sea level.

The river is in danger of flooding because of abnormally large amounts of rainfall. The National Weather Service issues a report stating that the river will rise 1 foot every 2 hours from where it is at two feet above normal at noon on Tuesday. They expect the crest to be at 20 feet. After the crest, they predict the river will fall at the same rate it rose.

Tom knows that the tide effect and the river crest usually add together.

Explicitness & Pictured:

Verbatim, likely: How high is the river crest (without tide effects) predicted to be?

Verbatim, unlikely: Who reports the rate that the river will rise?

Inference, likely: If the river bank can contain 25 feet at Tom's boat dock, should he worry about flooding?

Inference, unlikely: What would the water at Tom's dock taste like?

Table 2: Experiment 1 percent accuracy by task.

Task	Spatial		Non-spatial		Verbatim		Inference		likely	unlikely
	likely	unlikely	likely	unlikely	likely	unlikely	likely	unlikely		
overall										
summarize	75	77	70	84	88	82	57	79	76	
draw	88	82	80	78	91	85	59	80	80	
read-only	86	86	87	85	90	90	68	94	86	

Table 3: Experiment 2 percent accuracy by task.

Task	Spatial		Non-spatial		Verbatim		Inference		likely	unlikely
	likely	unlikely	likely	unlikely	likely	unlikely	likely	unlikely		
overall										
Immediate										
Summarize	75	83	82	82	85	88	66	84	81	
draw 90	86	87	89	84	84	58	87	83		
read-only	92	88	94	85	87	91	72	88	87	
Delayed										
summarize	75	74	79	80	86	75	56	78	75	
draw 82	83	84	83	87	85	56	90	81		
read-only	84	85	83	84	85	81	57	83	80	

Table 4: Experiment 3 percent accuracy by condition.

Condition unpict'd	Spatial		Non-spatial		Verbatim		Inference		pict'd
	Verbatim pict'd overall		Inference unpict'd		pict'd	unpict'd	pict'd	unpict'd	
Immediate									
illustrated	84	87	91	71	85	83	65	76	80
unillustrated	88	89	79	70	86	83	59	75	78
Delayed									
illustrated	83	82	87	82	83	70	59	73	77
unillustrated	84	85	72	80	79	77	50	71	75

Table 5: Experiment 4 percent accuracy by task.

Task	Spatial		Non-spatial		Verbatim		Inference		likely	unlikely
	likely	unlikely	likely	unlikely	likely	unlikely	likely	unlikely		
overall										
summarize	67	63	68	76	77	75	51	72	70	
draw	75	65	75	77	75	47	68	70		
read-only	79	78	82	85	83	75	57	75	77	
read-illus.	82	73	89	78	87	73	64	77	78	

FIGURE CAPTIONS

Figure 1. Experiment 1: accuracy by task. Error bars are \pm one within-group standard error from analysis treating subjects as the random factor.

Figure 2. Experiment 2: results by task and retention interval. Error bars are \pm one within-group standard error from analysis treating subjects as the random factor.

Figure 3. Experiment 4: accuracy by task. Error bars are \pm one within-group standard error from analysis treating subjects as the random factor.

Figure 1.

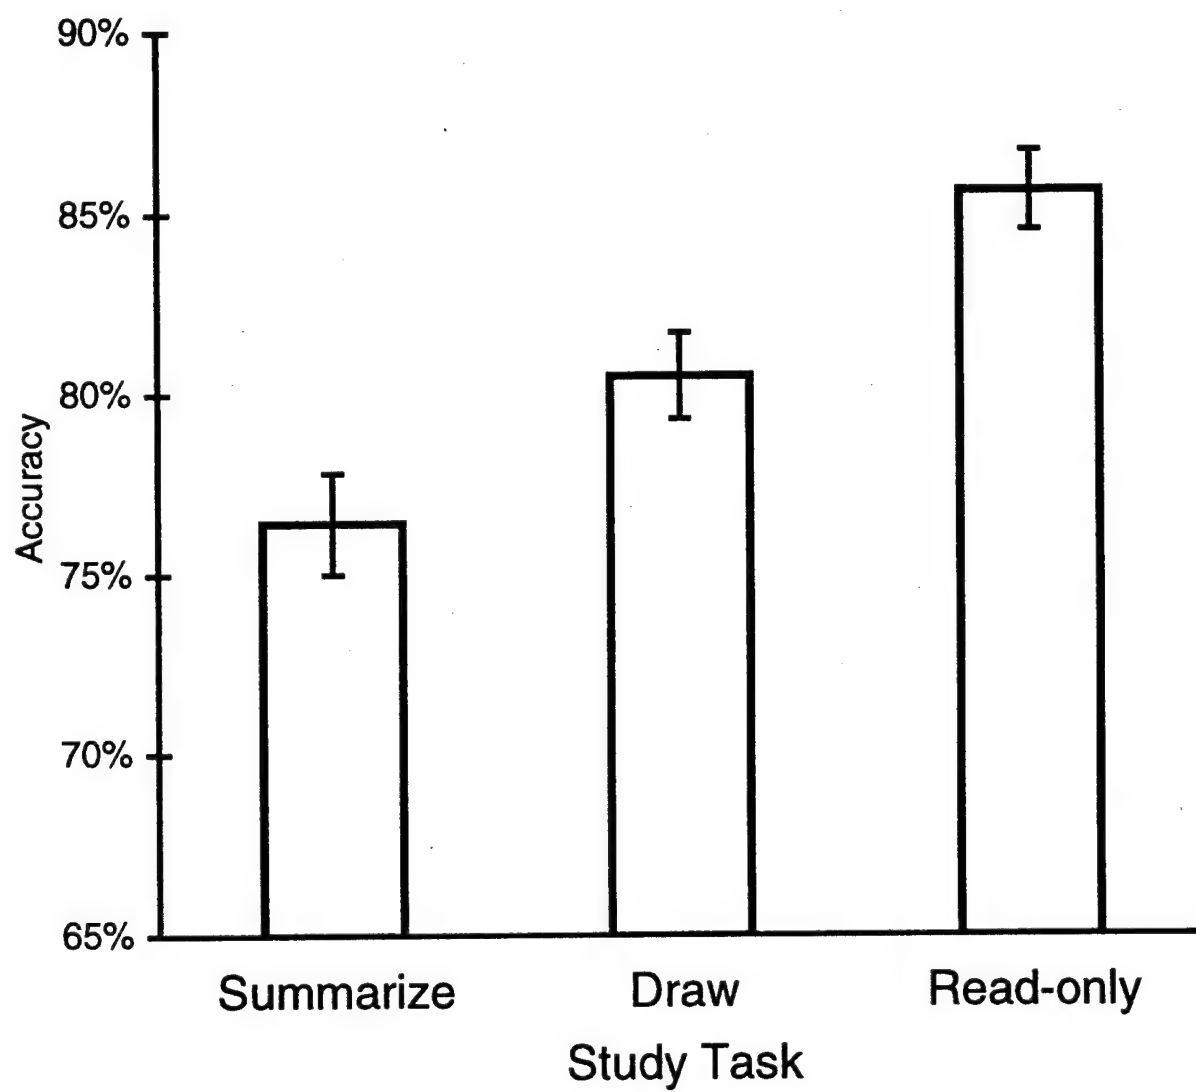


Figure 2.

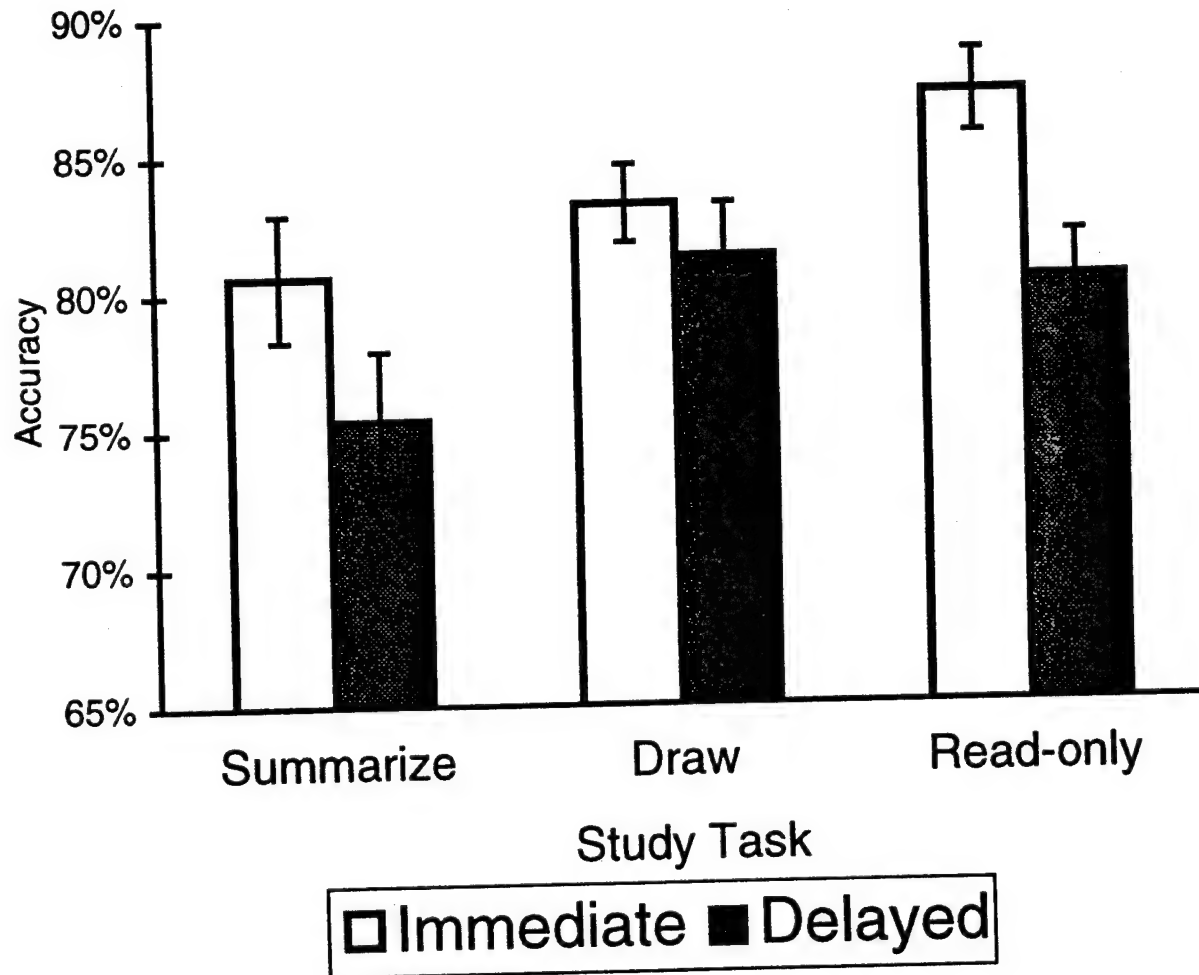
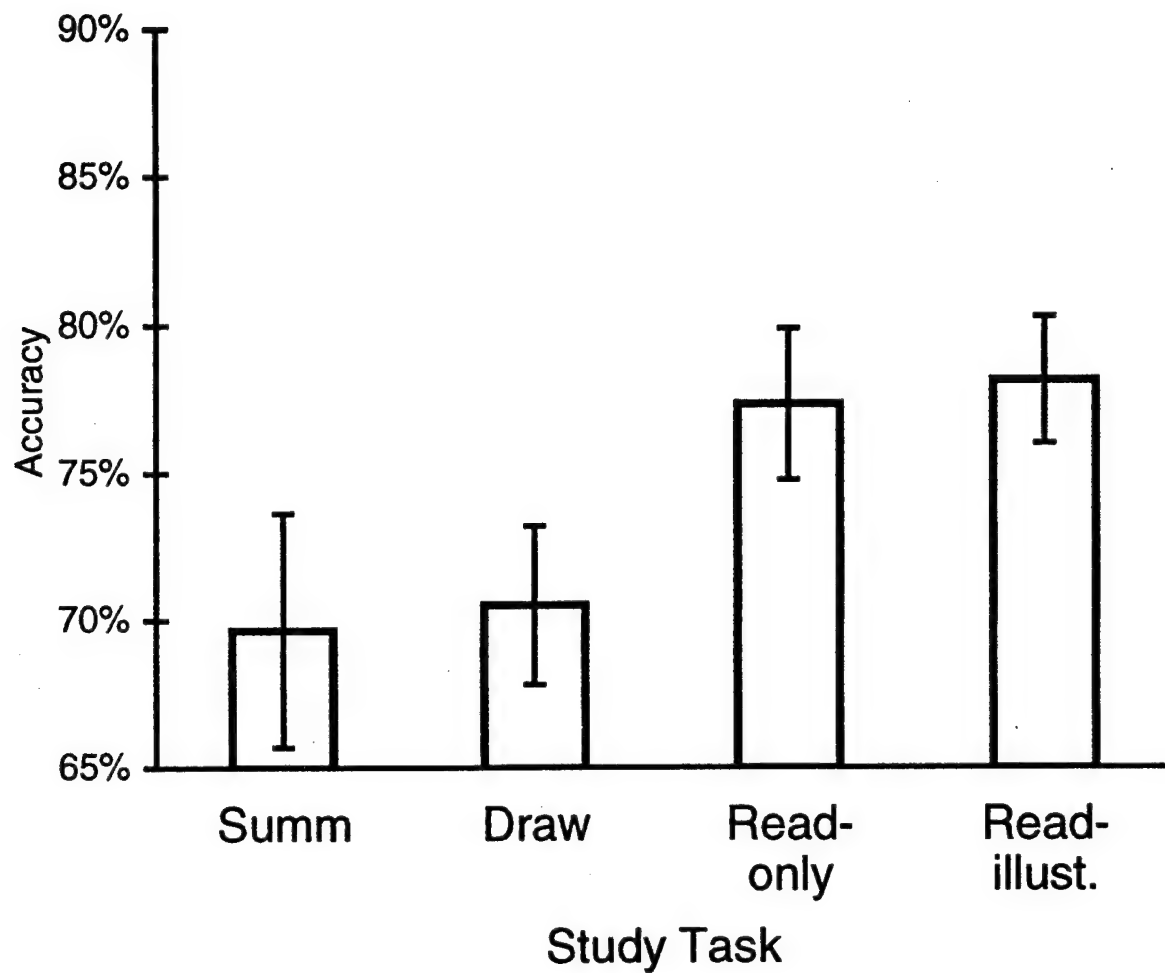


Figure 3.



Glenberg, full report

What memory is for

1. Introduction

Most memory theories presuppose that memory is for memorizing. What would memory theory be like if this presupposition were discarded? Here, I approach memory theory guided by the question, "What is memory for?" The answer that I develop is influenced by three sources. The first is Lakoff and Johnson's (Johnson, 1987; Lakoff, 1987; Lakoff & Johnson, 1980) cognitive linguistic analysis of language, conceptualization, and meaning. They propose that cognitive structures are embodied; that they arise from bodily interactions with the world (cf. Harnad, 1990, 1993). After a brief review of the Lakoff and Johnson program, I examine the literature on memory (the second source) for evidence that cognitive structures are, indeed, embodied, and why that is so. I will propose that memory evolved in service of perception and action in a three-dimensional environment, and that memory is embodied to facilitate interaction with the environment. The third set of ideas comes from research on mental model theory of language comprehension. I relate mental model theory to the notion of embodied memory by proposing that because language acts as a surrogate for more direct interaction with the environment, language comprehension must also result in embodied representations, which are in fact mental models. In exploring these ideas, I develop an approach to memory and language comprehension that suggests ways of dealing with old problems (e.g., why recollection and comprehension are effortful), as well as new concepts to replace old ideas (e.g., an association).

1.1 Why embodiment should matter to cognitive psychologists

Why should psychologists interested in language, learning, and memory care about issues such as embodiment of memory? Because by ignoring them, we have been making a big mistake. Most theories of memory treat internal representations as meaningless symbols such as a string of zeros and ones that "encode" features (e.g., Hintzman, 1986; McClelland and Rumelhart, 1986; Metcalfe, 1993), as point-like objects with no structure (Gillund and Shiffrin, 1984), or as propositions relating intrinsically meaningless symbols (Kintsch, 1988). Two problems arise from this treatment. The first is the symbol grounding problem (Harnad, 1990): How do those meaningless symbols come to take on meaning? The answer is not as simple as referring the symbol to a lexicon, because words in the lexicon must also be grounded. Also, not all of those meaningless symbols are meant to represent words or word-like concepts, but some are meant to represent complex non-verbal displays (Posner & Keele, 1968; Schacter, Cooper, & Delaney, 1990). The second problem is that we have not availed ourselves of a golden opportunity. By treating internal representation as meaningless symbols, we have not thought about the possibility of taking advantage of other forms of representation. Instead of meaningless symbols, suppose that representations have a structure that is lawfully related to the objects being represented. The structure of the representations might then play an important role in determining, for example, what concepts are easily associated because their structures literally fit together. For example, it seems easy to associate "horse" and "spotted" because horses have surfaces that can be spotted, whereas it is more difficult to associate "idea" and "spotted." Note that this sort of thinking trades on the analogical nature of the representations rather than on propositional listings of content (see Palmer, 1978). That is, we could just as easily assert "the idea was spotted" as "the horse was spotted." Nonetheless, one seems to make sense, the other does not.

In the next few sections I develop the case that internal representations are analogically structured (embodied), that this structure helps to explain memory phenomena, and that in conjunction these ideas suggest that the standard memory paradigms are ill-conceived and that standard memory phenomena may be revealing little that is important about memory. These

sections are followed by a discussion the possibilities for analogical representation underlying language comprehension.

1.2 Embodiment and The Lakoff and Johnson Program

A central concern of the Lakoff and Johnson program is the concept of meaning. According to Lakoff (1987), the standard theory of meaning in cognitive science is based on the notion of truth values of propositions, and as it turns out, this theory will not work as a theory of human meaning. Explication of why this is so requires a bit of patience, in part because the way psychologists use the term "proposition" is different from the way philosophers and logicians use it. For the psychologist, propositions are relations among symbols, that is, an assertion that a relation exists. It is these assertions that are supposed to be meaningful. Importantly, although the propositions are supposed to capture meaning, the symbols used in the propositions are taken to be, by themselves, meaningless or arbitrary: There is no intrinsic relation between a particular symbol and its meaning. Thus when illustrating propositions, a psychologist may use a word to stand for an element in the proposition, but that is just a convenience. Indeed, the meanings of the words need to be specified, presumably by other propositions. Thus, we should replace any words in a psychologist's proposition with things such as "symbol X19." This state of affairs is quite useful because it allows for reasoning (the derivation of new propositions) to be based on the manipulation of propositions by syntactic rules. These rules are thought to operate independent of the referents of the elements (nodes and symbols) in the propositions.

For example, suppose that proposition 1 (P1) asserts that element *a* is in relation *R* to element *b*. In shorthand, P1: *aRb*. Furthermore, suppose P2: *bRc*. Now, if *R* is a transitive relation (such as "larger than"), and both P1 and P2 are true, then by the syntactic rules of transitive inference, P3: *aRc* is also true. Thus, for the psychologist, we have created new knowledge, namely, that P3 is true. Note that these propositions have truth values, but they fail a common-sense test of what it means to have meaning. Namely, in order for a statement to be meaningful (to us), we must know what the statement is about. In contrast, although we know that P3 is true, we have little idea what it is about, because we have no idea what *a* and *b* stand for.

The problem of what *a* and *b* stand for is the symbol grounding problem (Harnad, 1990): How do we give meaning to the arbitrary symbols? To know what these propositions are about requires a mapping between the elements in the propositions (*a*, *b*, *c*, and *R*) and the world (or a possible world, or a model of the world). Without this mapping, the symbols can only refer to other symbols which in turn refer to yet other symbols. Just like trying to learn the meaning of a word in a completely foreign language by using a dictionary written solely in that language, such a system of symbols will never generate meaning (Searle, 1980). Most psychologists don't see a problem here, because they are happy to point to perception: The arbitrary symbols are grounded by the perceptual system. That is, what a symbol means is what it refers to in the "outside" world.

Lakoff (1987) presents (at least) three arguments against the plausibility of generating meaning by this sort of symbol grounding. First, this theory requires that categories be Aristotelian, that they have sharp boundaries. Aristotelian categories are needed so that we can successfully map between the arbitrary symbols in the propositions and the elements in the world. Thus, if a proposition is supposed to be about a horse, to give the proposition its proper meaning we must be able to map the symbol for horse (X19, perhaps) onto horses, and exclude zebras and antelopes and even perhaps ponies. In contrast to this requirement, there is a tremendous amount of empirical work in the psychology of human categorization implying that categories in the head are not Aristotelian. Instead many significant categories have fuzzy

boundaries (Oden, 1984, 1987), graded membership (Kalish, 1995), complex structures (Lakoff, 1987), or are based on prototypes (Rosch, 1973). Furthermore, the extensions of even basic biological categories are less than certain, and categories based on human culture are even more fuzzy. Thus, categories such as democracy, justice, and mother (Lakoff discusses biological mothers, birth mothers, adoptive mothers, step-mothers, etc.) seem to have structures quite different from the classical Aristotelian category.

A second argument against the standard theory as a theory of human meaning is based on an analysis of Putnam (1981). This analysis, however, is directed toward the philosopher's meaning of proposition, and so it requires a bit of new terminology. To the psychologist, a proposition (like aRb) is supposed to have meaning. To the philosopher, aRb is a sentence in a formal language. The meaning of the sentence (its propositional content) corresponds to the function that determines, for any possible situation, whether that sentence is true or false. In plain language, which is not an exact equivalent but close enough, the meaning of a sentence such as "the horse is spotted" is whatever allows one to determine if it applies to particular situations. Furthermore, two sentences have the same meaning if they have the same truth values for all possible situations.

Putnam discovered a serious problem with this truth-value notion of meaning: It is not difficult to construct pairs of formal sentences whose symbols are mapped to radically different things, but that have the same truth values in all situations. In other words, even though the sentences are about radically different things, on the truth-value account of meaning, because the sentences have the same truth values they are supposed to have the same meaning. Clearly, it does not make much sense to assert that sentences about different things mean the same thing. As it turns out, the problem is with the arbitrary nature of the symbols. They only mean when they are mapped onto the world, and Putnam demonstrated that it is impossible to find the one and only correct mapping.

Lakoff and Johnson's third argument against the standard theory is based on their analysis of language use and what it implies about cognition. In brief, people frequently use metaphorical language ("He's trapped in his marriage," "Your theory is airtight," "I'm really high today"). Furthermore, Lakoff and Johnson propose that metaphorical language is not just the way people talk, but that it accurately reflects the way people think. Given that theories cannot literally be airtight and that people's emotional states cannot literally be high, it is hard to imagine how cognition could be based on the mappings of arbitrary symbols and produce such (easy to understand) language.

Several other cogent arguments against the use of arbitrary symbols in a theory of meaning can be found in Barsalou (1993, Barsalou et al., 1993) and Shannon (1988). Barsalou and Shannon note that people have a hard time defining many familiar words, and that the definitions can vary greatly with context. This finding is difficult to explain if one believes that meaning of words is a simple list of well-formed propositions. They also note that there is no good account of how propositions composed of meaningless and arbitrary symbols might have evolved or how a child could have discovered them. As Shannon concludes,

"Specifically, it appears that the underlying substrate of mental activity is not a repertory of well-defined, well-structured abstract symbols, and that the workings of mind cannot be generally characterized as the computational manipulations of such symbols. Rather, the substrate in which mental activity takes place should be one that meets the following requirement: It should not be fixed by any coding system that is defined a priori, it should afford maximal sensitivity to unspecified dimensions and distinctions, it should be context-sensitive, and it should be embedded in the framework of

the organism's action in the world" (page 80).
That is a call for an embodied approach to meaning.

1.3 Embodiment and meaning

If we dismiss the standard theory, what is left? Lakoff and Johnson offer a theory of meaning based on the concept of embodied knowledge. Because I will be approaching the problem from the question of "What is memory for," I will develop an idea of embodied meaning that is distinct from the Lakoff and Johnson proposal. Nonetheless, the proposals are clearly related. In outline, my proposal is that perceptual systems have evolved to facilitate our interactions with a real, three-dimensional world. To do this, the world is conceptualized (in part) as patterns of possible bodily interactions, that is, how we can move our hands and fingers, our legs and bodies, our eyes and ears, to deal with the world that presents itself. That is, to a particular person, the meaning of an object, event, or sentence is what that person can do with the object, event, or sentence.

How does this approach answer the objections raised to the standard theory of meaning? Importantly, embodied representations do not need to be mapped onto the world to become meaningful because they arise from the world. In other words, embodied representations are directly grounded by virtue of being lawfully and analogically related to properties of the world and how those properties are transduced by perceptual-action systems (Harnad, 1990, 1993). Thus, the meaningful, action-oriented, component of conceptualization is not abstract and amodal. It reflects how bodies of our sort can interact with objects.

Given that embodied representations do not need to be mapped onto the world to be grounded, there is no need for representations to be Aristotelian nor for the categories in the world to be Aristotelian. Furthermore, because embodied representations are not discrete, meaningless symbols, they can reflect subtle, fuzzy variations in the world. How then do categories arise? Objects fall into the same (basic) category because they can be used to accomplish the same interactive goal, such as supporting the body. Because the same object may be useful for accomplishing a variety of goals, categorization can be flexible and context dependent (Barsalou, 1993).

Consider three objections to these claims. The first is that because we have different bodies, we will understand the world in different ways. In fact, that is a valid prediction. For example, what makes an object a chair for a particular individual will depend on whether or not that individual is able to get his or her body into a sitting position using the object. Thus, depending on the height of the object, the width of the flat surface, the object's strength, etc., the object will be a chair for some people (e.g., a child) but not for others (e.g., an aging grandfather).¹ Nonetheless, our bodies are substantially the same around the world and across cultures. Thus, although there will be variability around the edges, our common human endowments and our common environment ensures a great degree of common center to

¹ The argument is not that grandfather is unable to interpret a platform as a suitable chair for his granddaughter. Instead, the argument is that the grandfather knows perfectly well that the platform is not a suitable chair for himself. Thus, when the young child refers to the narrow platform as a chair, the grandfather may have some difficulty in knowing what the child is talking about. It is only when the grandfather adopts the child's perspective (that is, how the child's body can interact with the world) does he come to appreciate how the platform can be a chair for the child. (Adopting another's perspective requires suppression, as discussed in Section 3.4.) The main point is that words and events do not have set meanings for all times, places, and people. Words and events have meanings to individuals because of their individual bodies and experiences in the world.

cognitive structure.

A second objection to the claim that cognitive representation is embodied is that the mapping problem has not been solved; there is still the problem of mapping (arbitrary) words to embodied representations so that we can talk about what we are perceiving and thinking. This is a deep problem (e.g., Harnad, 1990; Plunkett, Sinha, Møller, & Strandsby, 1992), but it is not one that I intend to address here. The point of the above is that embodied representations allow us to understand how, except for the seriously deranged, we all know the difference between say, horses and ideas, and contrary to what Putnam's analysis shows of the standard theory, we don't ever confuse them.

A third objection is that some things are meaningful (e.g., a beautiful sunset) even when there is no apparent possibility for bodily interaction.² The embodied account of meaning is situated, so that action-oriented meaning can vary greatly with context. Thus, depending on the context, a coke bottle can be used to quench thirst, as a weapon, a doorstop, or a vase. That is, its meaning depends on the context. Similarly, a beautiful sunset is a context that combines with objects and memories to suggest actions consistent with warmth, relaxation, and a good beer.

Later I will discuss how embodied representations can be extended to represent abstract concepts and how they may provide a novel way of dealing with unanalyzed concepts such as association. For now, however, I turn to developing a particular sketch of embodied representations that arises from a consideration of what memory is for. This sketch is not a fully testable theory. The idea is to show how a type of theory that is not subject to the criticisms leveled at meaningless symbol theory can handle problems of memory and comprehension.

2. What Memory is For

Except for the recent blossoming of interest in indirect memory (see section 5.1), the contemporary psychology of memory has been dominated by the study of memorization.³ In part, this seems to have arisen from a failure of many twentieth century memory theorists to consider what memory is for. By the end of section 5 I will have concluded the following. Memory is embodied by encoding *meshed* (i.e., integrated by virtue of their analogical shapes) sets of patterns of action. How the patterns combine is constrained by how our bodies work. A meshed set of patterns corresponds to a conceptualization. Updating memory occurs whenever the meshed patterns change (a change in conceptualization of the environment), and the updating is in terms of a change, or movement, or *trajectory* toward a new set of meshed patterns. Thus, memory records how conceptualizations blend into one another. This memory works in two broad modes. First, patterns of action based on the environment (projectable properties of the environment) are automatically, that is, without intention, meshed with patterns based on previous experience. This automatic use of memory corresponds closely to implicit or indirect memory. Second, patterns from the environment can be suppressed so that conceptualization is guided by previous experience encoded as trajectories. This is a conscious and effortful use of memory. The ability to suppress environmental patterns contributes to prediction, the experience of remembering, and language comprehension.

2.1 The function of memory in a dangerous environment

We live in a dangerous, three-dimensional world. Given the size, density, and physical capabilities of our bodies, the natural environment is hostile. We are open to predation, and our

² Thanks to Alan Baddeley for raising this objection most forcefully.

³ Thanks to Tony Sanford for this observation.

interactions with the world can lead to injury from freezing, burning, drowning, and falling. Clearly, survival requires the capability to navigate this environment, and just as clearly, our perceptual system has evolved to do just that. For example, we have developed impressive abilities to use information (e.g., optical flow fields) to guide action so that obstacles are avoided. These abilities may not require any sort of representation of the environment, and they may not require memory; responding constrained by characteristics of the environment and our bodies guarantees successful action (for a review, see Bruce & Green, 1985).

On the other hand, it is frequently the case that we need to differentiate. In addition to avoiding obstacles in our path, we need to pick out and follow a particular path, avoid a particular location, or approach a particular person. This sort of differentiation requires a memory system. What makes one person a particular person (to you) or one path the path to your house, is its relevance to you, that is, how you have interacted with it in the past. An optical flow field cannot contain this information; it is the province of memory. This distinction is discussed by Epstein (1993) who uses the term "projectable" to refer to properties of the environment that can be specified by information available in the light and "nonprojectable" to refer to properties that must be signaled by other sources. Thus spatial layout is a projectable property whereas ownership is a non-projectable property that must come from experience.

2.2 Embodied conceptualization, memory, and meaning

To support action, the perception of projectable properties is in terms of patterns of possible action: how we can examine, grasp, shove, leap over, or move around an object. This coding depends on the capabilities of our bodies, both as a species and as individuals. Because the world is perceived in terms of its potential for interaction with an individual's body, it is proper to call the perception "embodied."

Patterns of action derived from the projectable properties of the environment are combined (or meshed, section 3.1) with patterns of interaction based on memory. The two patterns can combine because they are both embodied, that is, both are constrained by how one's body can move itself and manipulate objects. The resulting pattern of possible actions is a conceptualization: the possible actions for that person in that situation. For example, turn left to get home.

Thus meaning of an object or a situation is a pattern of possible action. It is determined by the projectable features of the object molded by bodily constraints and modified by memory of previous actions. These memories provide the non-projectable features. As another example, consider the meaning of the cup on my desk. the embodied meaning is in terms of how far it is from me (what I have to do to reach it), the orientation of the handle and its shape (what I have to do to get my fingers into it), characteristics of its size and material (the force I must exert to lift it), etc. Furthermore, the meaning of the cup is fleshed out by memories of my previous interactions with it: pouring in coffee and drinking from it. Those memories make the cup mine.

Note three characteristics of this sort of meaning. First, because bodily actions take place in space, embodied meaning captures spatial (or topological) and functional properties. Thus a synonym for this type of embodied meaning is spatial-functional meaning. Second, because we interact with objects via parts, conceptualization in terms of bodily interaction forms the basis for partonomies (Tversky and Hemenway, 1984) and basic-level categorization. Third, conceptualization in terms of patterns of bodily interaction is very close to Gibson's (1979) notion of affordance.

Thus, what is memory for? Its primary function is to mesh the embodied conceptualization of projectable properties of the environment (e.g., a path or a cup) with

embodied experiences that provide nonprojectable properties. Thus the path becomes the path *home* and the cup becomes *my* cup. This meshed conceptualization, the meaning, is in the service of control of action in a three-dimensional environment.

2.3 Evidence for embodied conceptualization

How far can this account of embodied meaning be pushed? At the least, there are intriguing results that fit this account nicely and which do not seem to have a natural explanation in cognitive accounts based on meaningless symbols. I will review some of this literature from domains of affect, memory, and imagery.

2.3.1 Embodiment and affect

Van den Bergh, Vrana, and Eelen (1990) presented typists and non-typists sets of letter pairs (e.g., WX and ZD). The subjects were asked to choose the one pair (from each set) that was liked the best. Typists showed a clear preference for pairs typed with different fingers over pairs typed with the same finger, whereas the non-typists showed little preference. (The typing finger was determined using AZERTY keyboards in Belgium and QWERTY keyboards in the U.S.) Van den Bergh et al. argued that for typists, part of the encoding of letters is as a motor program or movement. The incompatible movements generated by letters typed with the same finger resulted in a negative evaluation. It is unlikely that this effect arose from associations to specific letter combinations because the effect was most robust for pairs of letters with low frequency in the language.

Berkowitz and Trocolli (1990) and Berkowitz, Jo, and Trocolli (1993) illustrate the influence of the body on affect judgments. In one experiment, subjects were asked to judge the personality of a fictitious person described in neutral terms. Half the subjects listened to the description while holding a pen between their teeth without using their lips. This activity forces the face into a pattern similar to that produced by smiling. The other subjects listened to the description while biting down hard on a towel. This activity forces the face into a pattern similar to that produced by frowning. The subjects who were smiling rated the person described more positively than did the subjects who were frowning. It is unlikely that this effect arose due to demand characteristics of the experiment for the following reason. The effect was obtained only when subjects were distracted from their activities; when they were asked to focus on the activities, the subjects seemed to compensate for the forced smile (frown) and rate the description more negatively (positively). What can account for this finding? Experienced emotion is embodied. When the body is manipulated into a state that is highly correlated with an emotion, the body constrains other cognitive (that is embodied) processing.

2.3.2 Embodiment and imagery

Montello and Presson (1993) asked subjects to memorize the locations of objects in a room. The subjects were then blindfolded and asked to point to the objects. Pointing was fast and accurate. Half of the subjects were then asked to *imagine* rotating 90 degrees and to point to the objects again. That is, if an object was originally directly in front of the subject and the subject imagined rotating 90 degrees clockwise, the correct response would be to point to a location toward the subject's left. In this condition, the subjects were slow and inaccurate. The other subjects, while blindfolded, were asked to *actually* rotate 90 degrees and to point to the objects. These subjects were just about as fast and accurate as when pointing originally. Thus, mentally keeping track of the locations of objects, a task that many cognitive psychologists would suspect as being cognitive and divorced from the body, is in fact strongly affected by literal body movements.

Rieser, Garing, and Young (1994) reported a similar finding for children and adults. The participants were tested for the ability to imagine (while at home) their classrooms and to point to objects from various perspectives. When the perspective change was accomplished by actually changing position (at home), the 5-year-olds were correct on 100% of the trials, the 9-year-olds were correct on 98%, and the adults on 100%. When the perspective change was accomplished solely by imagination, the 5-year-olds were correct on 2% of the trials, the 9-year-olds were correct on 27%, and the adults were correct on 100%. Even the adults showed great difficulty in terms of the time needed to accomplish the imagination-only version of the task. When actually changing position, 100% of the adult responses required less than two seconds, whereas when only imagination was used, only 29% of the responses required less than two seconds.

Findings on the psychophysiology of imagery also point to a close connection between body and cognition. These findings are summarized by Cuthbert, Vrana, and Bradley (1991). Their starting point is Lang's (1979) bio-informational theory, which asserts that encoding of events includes response "propositions", and that imagery (visual and otherwise) is the activation of those propositions. Furthermore, although overt responding is inhibited during an imagery task, there may well be "efferent leakage" that can be measured using psychophysiological techniques. In support of these ideas, Cuthbert et al. note that psychophysiological responsivity is particular to the image being evoked. Thus imagining a fearful situation evokes sweating, imagining positive situations results in measurable activity in muscles associated with smiling, and imagining negative situations results in activity in the muscles associated with furrowing of the brow. There are analogous effects for imagery related to other perceptual/action systems. Thus, in imagining pendular motion, discharges in the eye muscles follow the appropriate frequency, in imagining bicep curls there are discharges in the biceps, and in imagining the taste of a favorite food there is an increase in saliva flow. These results are compatible with the notion of embodied, spatial-functional encoding. In addition, the idea of embodied encoding has an advantage over Lang's response propositions. According to Cuthbert et al. (1991), the function of imagery is to allow new behaviors to be tried out "off-line." It is not clear however, how response propositions can be integrated (other than by concatenation) to effect this rehearsal. In contrast, the integration of responses is basic to the notion of mesh (see Section 3.1) of embodied encodings. That is, given that the information is encoded in terms of bodily interaction, effecting one action (or imagining it) necessarily constrains the operation of simultaneous and successive actions.

2.3.3 Embodiment and memory

Effects of embodiment is revealed by research on memory for subject-performed tasks (Cohen, 1981; Engelkamp & Krumnacker, 1980; Saltz & Donnenwerth-Nolan, 1981; see also a special issue of *Psychological Research*, 1989). The basic finding is that memory for actions (performing a command such as, "open the book") is better than memory for the verbal description of the commands. One interpretation of this finding is that memory specializes in embodied information.

The nature of our bodies also controls ease of remembering. Consider a series of studies by Tversky and her colleagues (e.g., Bryant, Tversky, and Franklin, 1992). In these experiments, subjects read about and memorized spatial layouts corresponding to scenes viewed from particular perspectives (e.g, in the hotel scene, "To your left...you see a shimmering indoor fountain..."). Objects were located above, below, in front, in back, to the left, and to the right of the observer in the imagined scene. After the scene was memorized, the time taken to retrieve a particular object was measured. For equally well memorized locations, one might expect the retrieval times to be independent of location. Another hypothesis is that the times would be correlated with the degree of mental rotation needed to mentally face the object. The

results, however, were contrary to both of these hypotheses. Fastest responding was to objects located on the head/feet axis, followed by the front/back axis, followed by the left/right axis. Tversky argues that these results follow from using a "spatial framework" that is sensitive to environmental asymmetries (such as gravity) and perceptual asymmetries (we generally look and attend to the front). In other words, retrieval processes appear to be sensitive to how we use our bodies.

Klatzky, Pellegrino, McCloskey, and Doherty (1989) demonstrated contributions of the body to symbolic or semantic judgments. They trained subjects to make hand shapes corresponding to descriptions such as "pinch" or "clench." The verbal descriptions were then used as primes for judging the sensibility of phrases such as "aim a dart" (sensible) or "close a nail" (not sensible). The appropriate prime, that is a prime corresponding to the hand shape used in the to-be-judged action, speeded the sensibility judgment compared to a neutral prime. Thus, the hand shape for "pinch" speeded the sensibility judgment for "aim a dart." It is unlikely that this priming effect derives from any sort of verbal mediation: The priming effect was found for subjects trained to make the hand shapes when signaled by non-verbal primes. Also, when subjects were trained to make verbal responses (but not hand shapes) to the non-verbal primes (e.g., saying the word "pinch" when shown the non-verbal signal for pinch), the priming effect was eliminated. Klatzky et al. suggest that the sensibility judgment requires a type of mental simulation using an embodied, motoric, medium. Generating the appropriate hand shape "facilitates constructing the representation and/or simulating the action/object pairing."

3. How embodied memories are used

Consider this scenario. You have been wandering in the woods, and suddenly you are unsure of the way home. You see what appears to be a path, but you are not certain if it really is a path, yet alone, the path home. You take a few steps and hunt for evidence. As you continue your exploration, you become convinced that this is the right path: the patterns of rocks, twigs, and soil align themselves to form a connected pattern that could be a path. Also, as you move along, you are able to conform your own body to the putative path. For example, the overhanging branches are not so low that you have to stoop or crawl; when you reach a stream, the distance between the rocks forms a series of stepping stones that can be used by an animal of your size and agility.

3.1 Mesh of patterns as functional constraint satisfaction

Recognition of the path as a path arises from an exploration of the environment and a fit between the environment and embodied knowledge. This fit can be conceptualized as a type of constraint satisfaction, but here the constraints are spatial and functional, not associationistic or probabilistic (cf., Rumelhart, Smolensky, McClelland, & Hinton, 1986). Thus projectable properties of the environment (arrangement of rocks, twigs, and soil) are encoded in terms of how you (with your particular body) can interact with that environment (e.g., whether the distances between the rocks in the creek can be broached). Other patterns of interaction come from memory, for example, patterns representing goals such as "get home without getting wet." In conceptualizing the environment as a path, the spatial-functional patterns based on projectable properties from the environment are combined or meshed with the patterns from memory. The meshed pattern dictates how (or if) the body can be moved in a way that simultaneously satisfies both sets of patterns of action (e.g., "Can I, with my body, get from rock to rock without getting wet?"). This sort of mesh is a possibility because all of the patterns are embodied, that is, they are all encoded in terms of how your body constrains actions. When the patterns can be meshed into a plan for coherent action (e.g., stepping across the rocks), the rocks, soil, and twigs become (for you) a path.

I envision mesh of embodied encodings as being analogous to coarticulation in speech production. When pronouncing the initial /d/ in "dog," the articulators are shaped in part by the requirement to enunciate the following vowel, and when pronouncing the vowel, the articulators are shaped not only by the vowel, but by the preceding and following consonants. Furthermore, the constraints on articulation are not consciously imposed, but are constraints that follow from real movements of physical devices: The tongue can only be in one place at one time, and how it is going to move to the next place will depend on where it is now. Thus pronunciation of the word requires a mesh of real physical actions.

An example of cognitive meshing is borrowed from Barsalou, Yeh, Luka, Olseth, Mix, & Wu (1993). Imagine a ball; now imagine that it has yellow and white stripes; now imagine that the ball is deflated (it is a beach ball). Adding each new descriptor is not a matter of adding a simple association or adding a proposition to a list. Instead, each previously constructed representation constrains how the new descriptor is utilized. Thus, the yellow and white stripes surround the ball. Then, not just the ball, but the stripes too become deformed when the ball is deflated. The stripes and the ball deflate together because they are encoded as patterns of action subject to the same spatial-functional constraints. This meshing occurs, not just in imagination, but in memory, comprehension, and perception.

It is the mutual modification of meshed patterns of action that produces emergent and creative features of thought. The deflated beach ball is not simply a deflated ball associated with an unchanging stripes feature. Instead, the fact that the stripes are deflated arises from the operation of meshing. Related concepts will mesh easily, because that is what it means to be related (section 7.3), and with some effort, we can mesh arbitrary concepts. Thus a "tiger bicycle" is one designed for hunting tigers, and it consists of a mesh between the actions required to hunt and those required to ride a bicycle, whereas "colorless green ideas" are uninspired ways of dealing with environmental crises. In short, mesh underlies our ability to understand novel conceptual combinations. Note that the type of mesh I am proposing depends on the analog nature of embodied actions, not just their propositional content.

3.2 Clamping projectable properties

Meshing patterns of action based on projectable properties of the environment with those from memory changes the way we conceptualize the environment. Thus, the soil, twigs, and rocks are conceptualized not just as a path, but as the path home. There is a danger, however, in allowing patterns from memory to modify conceptualization: meshing of patterns can distort the perception of the environment. Clearly, survival requires seeing the environment for what it is (soil, twigs, rocks), not just for what it means (the path home). To keep the system reality-oriented, it is necessary to ensure that patterns based on projectable properties of the environment are primary. That is, the meshed conceptualization that is achieved cannot be at the cost of distorting the environmental input. I will refer to this as clamping projectable properties of the environment.

Clamping projectable properties ensures that experiences are individuated or situated. We do not experience categories, but individual, particular events (cf. Barsalou et al., 1993). We cannot direct our perceptual system to ignore differences between two paths, just because they are both paths, or between two chairs just because we can fit our bodies into both. Because the projectable properties are clamped, the two chairs, although members of the same category, remain separate chairs.

3.3 Updating memory

I have proposed that embodied memory acts as a source of nonprojectable patterns of

action that mesh with patterns derived from projectable properties of the environment; the mesh is possible because both sets of patterns are constrained by how the body works. If memory is to be useful, however, it must be updated. That is, new experiences must affect the system so that we come to learn the path home. Because experience is continuous (or at least the environment appears continuous to beings of our size and abilities), we must deal with how it can be captured by a system using a finite brain.

Consider this possibility. Projectable properties are clamped and then embodied memories mesh to produce a particular conceptualization (e.g., the path home). At this point, either an action is taken (e.g., a step along the path) or projectable properties of the environment change (e.g., a barrier appears). In either case, the system is forced to settle into a new conceptualization. Here is the proposal for updating memory: Memory is updated automatically (that is, without intention) whenever there is a change in conceptualization (mesh). The degree to which updating takes place is exactly correlated with the degree to which the conceptualization changes.

Updating is not encoding a new memory trace. Instead the shift from one pattern of possible actions (one conceptualization) to the next is reinforced. That is, what is updated is how one situation flows into another. I will refer to this flow as a trajectory. I am using the term trajectory to imply that the change is not random. Instead, actions humanly possible under the current conceptualization are biased by what was possible in the previous conceptualization, just as pronunciation of a vowel is biased by the pronunciation of the preceding consonant.

The idea of trajectories solves several problems in the psychology of memory. It provides a way of conceptualizing dynamic information in memory that is sensitive to biological and spatial-functional constraints (Shiffrar, Heptulla, O'Shaughnessy, & Freyd, 1993). Trajectories can reflect minimal changes in conceptualization, such as from one step along a path to another, or gross changes such as from a step to a fall. The idea offers the beginnings of a solution to the problem of features. Most theories of memory are based on the idea that memories are multidimensional, consisting of a vector of features, such as animate, red, and smaller than a bread box. None of these theories, however, is committed to a listing of what those basic features might be. In fact, because experience is so varied, it is hard to imagine a complete list. Also, given a feature-based system, it is difficult to understand how people can ever learn anything truly new: We must always conceptualize using the same basic features. In contrast, because embodied patterns of action can be infinitely varied and infinitely meshed with goals (also specified as patterns of action), a system based on embodied concepts and trajectories approaches the ideal of enabling memory to code the full variety of human experience.

Because updating of trajectories occurs only when there is a change in conceptualization, memory is sensitive to frequency and to novelty. To illustrate this, consider once again walking the path home. Three phenomena are associated with repeated actions. (1) Memory for the repeated action (walking the path) will be an increasing, but negatively accelerated, function of frequency (e.g., Logan, 1988). (2) Memory for a particular typical repetition of the action will be poor (Glenberg, Smith, & Green, 1977; Naveh-Benjamin & Jonides, 1984). (3) Memory for a particular unusual repetition of the action will be good (Hunt, 1995). The frequent interactions with the path will result in frequent updating (reinforcement of a particular trajectory), and consequently a shift toward a stable conceptualization (e.g., a shift from possibly the path home to definitely the path home). However, once the conceptualization is stable, little further updating occurs. Thus, each encounter with the path will have less and less of an impact (phenomenon 1). Because typical encounters result in little new conceptualization and little updating, we have little memory for the individual steps down the path (phenomenon 2). However, if reconceptualization is required (e.g., when a log appears across the path, so that now, in terms of bodily constraints on action, the path is a blocked path) memory is again

updated, leading to memory for novel events (phenomenon 3).

3.4 Prediction and suppression of the clamped environment

The meshed conceptualization of the current environment dictates what actions are possible in that environment. Prediction, however, requires simulating how an action will produce a new conceptualization, which in turn can be used to simulate the next action, etc. Two difficulties arise. The first is that simulated action does not change the environment. Thus, changes in projectable properties that would have resulted from a real action cannot be clamped to automatically guide further action. A second difficulty is that currently clamped stimulation is providing the *wrong* constraints, because those constraints are only relevant before the simulated initial action. I believe that this is a major problem, and that it requires a radical (and dangerous) mechanism: suppression. In particular, I propose that in the service of prediction, we have developed the ability to, if not ignore, at least to suppress the overriding contribution of the current environment to conceptualization. This is a risky operation because it loosens the tie between reality (the current environment) and conceptualization. Perhaps because suppression is so dangerous, it is an effortful process. As we will see, however, suppression results in several serendipitous abilities, including conscious autobiographical memory and language comprehension.⁴

Once clamping of projectable properties is suppressed, multi-step prediction arises from following trajectories guided by bodily constraints on action. For example, by following trajectories we can envision what will happen when we proceed down the path. We also have the ability to envision arbitrary events (such as what actions are possible if the path is washed out by a storm or blocked by strange creatures) not just events we have previously experienced. Prediction for these arbitrary scenarios is based on seeking a mesh among patterns of action. Some of the patterns are based on well-learned trajectories. Other patterns (e.g., interactions with strange creatures) come from a consideration of how our bodies work. These patterns can mesh to give a coherent conceptualization because they are all based on bodily interaction. Keep in mind, however, that in prediction, the mesh of these patterns may not be guided by stable and projectable features of the environment. To the extent that environmental constraints are suppressed, and to the extent that trajectories are not well-learned, the predictions will tend to be variable and inaccurate. Thus, it is easy to predict the outcome of the next step on a well-traveled path: the simulated mesh is strongly constrained by projectable features of the current environment and well-learned trajectories. It is more difficult to predict what will happen many steps down a new path when the projectable features must be suppressed and trajectories uncertain.

3.5 Mesh and connectionism

Many of the ideas and much of the terminology introduced in Section 3 are borrowed from connectionist approaches to cognition. Some examples are constraint satisfaction, trajectories as paths through a set of states, and clamping of projectable features. There are two other, perhaps deeper, similarities. As I will discuss in Section 6.2, an embodied conceptualization functions as a preparatory state. Given a particular conceptualization, an organism is better prepared to act when changes in the situation easily mesh with the conceptualization than when changes do not easily mesh (i.e., we are surprised). This notion of preparedness underlies priming phenomena, and it is close to connectionist interpretations of semantic priming developed by

⁴ Furthermore, Gärdenfors (1995) proposes that the ability to form "detached" representations, that is, those unconnected to current environmental stimulation, is a significant evolutionary step. Bloom (1994) discusses how the ability to manipulate such representations motivates language acquisition.

Masson (1995) and Sharkey and Sharkey (1992).

The second deeper similarity relates to ideas of context and situated representation. For example, Smolensky (1988) discusses how a distributed representation of "coffee" will depend on whether the coffee is in a cup, in a can, or in a person. Similarly, as I discuss in Sections 3.1 and 7.3, action patterns based on projectable features of an object (e.g., a coke bottle) can mesh with action patterns underlying goals in particular contexts (e.g., drinking or fighting), so that the resulting meshed conceptualization is context-dependent.

Nonetheless, there are important differences between my use of terminology and connectionist systems. For example, connectionist accounts of semantic or meaningful information are based on conceiving of meaning as activation of a limited number of features, at least at the input layer. Unfortunately, most theorists fail to specify what the features are, and they fail to specify how those features might be learned or changed as a consequence of development. In the system that I am proposing, initial coding is not featural, but analog, in terms of patterns of possible action. Furthermore, as one learns more about the interactive capabilities of one's body, objects and actions can be imbued with new meaning: What I can do with that object now.

A second important difference concerns the nature of constraints. In standard connectionist accounts, constraints are, in Palmer's (1978) terminology, extrinsic (but see Regier, 1995, for an exception). That is, a particular constraint represents statistical, or joint occurrence, information, not a necessary feature of the operation of the system. Thus, a connectionist system would be equally happy to learn that a coke bottle can be used as a chair or as a weapon. In an embodied system, constraints arise because of analog coding of projectable features and their implications for human action. In Palmer's terminology, these constraints are intrinsic to the operation of the system. For example, how we think about a coke bottle is constrained not just by particular experiences with coke bottles, but by the actual shape and heft of the bottle, too. Thus, an embodied system would have little difficulty understanding how a coke bottle could be a weapon, but it would balk at learning that it could be used as a chair.

These differences are not unique to my proposal. Lakoff (1988) argues that connectionist systems need to be grounded in the body to give meaning to connections and constraints. As an illustration, he notes that phonology is not arbitrary; instead, it is constrained by the muscles, shapes, and control of articulation. Shepard (1988) makes a related point regarding the abilities of connectionist systems to self-organize and generalize:

"But nontrivial self-programming can take place only if some knowledge about the world in which the system is to learn is already built in. Any system that is without structure has no basis for generalization to new situations." (p. 52).

How the body can interact with the world provides just such a basis for generalization.

These comments should not be taken to mean that an embodied system cannot be simulated using connectionism. In fact, it may well be that connectionism will be the surest route to formalizing these ideas. Nonetheless, it will have to be a connectionism that differs from the sorts currently in use.

4. Memory in the long-term and the short-term

The system described so far seems to be useful for negotiating the environment, and it seems to correspond to what some have called semantic (Tulving, 1983) or generic (Hintzman,

1986) memory. Where is episodic memory, that is, our memory for particular, personal experiences? The answer: the same place. I propose that episodic recollection is a type of pattern completion via meshed bodily constraints on action. Furthermore, the episodic character, the feeling that a memory is personally relevant, arises from suppressing clamped projectable properties of the environment. In this case, conceptualization is driven by trajectories rather than by changes in the environment.

To some cognitive psychologists, this idea will seem wrong on the face of it: It denies the difference between episodic and semantic memory; it denies the idea that episodic memory is temporally organized; it provides no distinction between short-term and long-term memory. Before describing how the idea seems right, I will briefly address why these problems are more apparent than real.

4.1 Episodic and semantic memory

I am explicitly equating episodic and semantic memory in the sense that there are no separate episodic and semantic memory systems, hierarchically arranged (Tulving, 1984) or otherwise. Of course, phenomenal memories differ in content, accessibility, etc. But those differences do not imply separate systems. Whereas this equation of memory systems may have been controversial ten years ago, data and mainstream memory theorizing are moving in this direction. In short, there is little data to support a distinction between a memory system devoted to personal experience and one devoted to general knowledge (McKoon, Ratcliff, and Dell, 1986). What appeared to be strong evidence for a memory organized by "semantic" dimensions (Collins & Quillian, 1969), is now known to reflect frequency of experience (Conrad, 1972). Evidence that was taken to indicate the storage of prototypes (Posner and Keele, 1968) in semantic memory, is now taken to reveal the operation of retrieval processes that can average experiences (Hintzman, 1986; McClelland and Rumelhart, 1986). Priming effects that were thought to reflect the spread of activation along permanent semantic links can be easily demonstrated for newly learned (hence episodic) information (McKoon and Ratcliff, 1986a). Thus the distinction between episodic and semantic memory probably reflects a difference in the frequency with which the memories are used, the methods of assessment, and the content of the information, rather than any intrinsic differences in memory systems.

4.2 Temporal organization of episodic memory

If the framework that I have described is the only memory system, then it explicitly denies a tenet of theorizing about episodic memory: Memory is a record of events that maintains some semblance of temporal order [see, for examples, Murdock's conveyor belt model (1974) or Glenberg and Swanson's (1986) temporal distinctiveness theory]. Almost assuredly, the tenet that episodic memory maintains order derives from the fact that temporally distant information is harder to remember than recent information. This fact does not demand a theoretical explanation that maintains time as a dimension of memory, however. In fact, Friedman (1993) presents a convincing case that episodic memory is not organized temporally. First, there is little priming between temporally contiguous but are otherwise unrelated experiences. Second, memory for time of occurrence of events is not just inaccurate, it shows nonmonotonic scale effects. That is, memory for when an event occurred may be accurate for the day, inaccurate for the month, accurate for the season of the year, inaccurate for the year, but accurate for the decade. Third, as Friedman discusses, for most of human history, memory based on a linear dimension of time would serve little useful purpose. Instead, a memory organized by functional significance, or by recurrent events (seasons, migrations, life cycles), would seem to have much greater adaptive significance.

4.3 Short-term memory

The idea of a single memory system seems wrong in that there is no mention of separate processes for long-term memory and for short-term or working memory. Much of the evidential basis for a separate short-term store (or working memory, according to Baddeley, 1990) has been eroded. For example, the recency effect is the enhanced recall of items from the end of a list. Because it was thought to be easily disrupted by a short period of distraction, it was taken as a hallmark of short-term store. We now know, however, that recency effects can be very long-term (Glenberg, 1984; Greene, 1986, 1992). Another supposed hallmark of a separate store is acoustic/articulatory encoding (e.g., Hintzman, 1968). That is, short-term store was believed to code information along acoustic/articulatory dimensions, whereas long-term store coded "semantic" information. However, demonstrations of meaning-like coding in short-term situations (Shulman, 1972), as well as long-term memory for articulatory and orthographic information (e.g., Hunt & Elliot, 1980) deny this simple distinction. Also, the quick forgetting demonstrated using the Brown-Peterson distractor technique, is now known to reflect a combination of poor initial coding (Muter, 1980), and interference from previously studied material (Keppel & Underwood, 1962; Watkins & Watkins, 1975).

What are we to make of the impressive body of information on apparently separate short-term modules (e.g., Baddeley's articulatory loop, phonological store, and visual/spatial sketchpad)? An alternative theoretical position is to consider the evidence as indicative of skills and strategies effective in particular domains (cf. Kolers & Roediger, 1984), rather than separate modules. This skill-based alternative can easily accommodate findings that might otherwise be interpreted as evidence for new working memory modules. As one example, Reisberg, Rappaport, and O'Shaughnessy (1984), demonstrated an increase in "working memory capacity" by instructing subjects how to use their fingers to code numbers in a memory span task. This evidence might be interpreted as evidence for a new "finger-control" module, but it seems more sensible to view it as a newly-learned skill. As another example, Carpenter, Miyake, and Just (1994) speculate that there may be separate working memory capacities for language production and language comprehension. Again, the alternative that different skills are involved in comprehension and production would seem to more easily accommodate the data.

Nonetheless, one must come to grips with our intuitions of immediate access to some information and difficulty in recovering other information. Consider this proposal. Memory and the perceptual/action system are designed to produce a meshed conceptualization (possible actions) for current stimulation. It is this constantly changing conceptualization (changing because the stimulation changes in response to action) that gives the illusion of a short-term memory. Because the current conceptualization updates memory and provides the starting point for future conceptualization, it will have a strong influence on performance over the next few moments (as does a short-term memory). Distraction (a changing environment) does cause a disruption in short-term behavior because it produces a forced change in the current conceptualization. Limits on the "capacity" of a short-term store are simply the limits on coherent conceptualization.

This framework also rationalizes some aspects of rehearsal and control of thought. In particular, it seems that some sort of cyclical activity is needed to maintain information in the forefront of consciousness. Baddeley (1990) discusses this as an articulatory loop that must reactivate the decaying contents of a phonological store. But if memory is like a box that holds items of information, why should cyclical activity be necessary? The answer comes from the nature of trajectories. They are not static memory traces; they are reinforced changes from one conceptualization to the next. Thus, there is no holding of trajectories in mind. Instead, to maintain a thought or a conceptualization in the absence of clamped projectable properties, it is necessary to reuse a trajectory, or to replay the same scene over and over.

5.0 Memory in two modes: automatic and effortful

The major function of memory is to mesh constraints on action based on non-projectable properties with constraints from projectable properties. This is an automatic function of memory in the sense that it is not under conscious control, and it corresponds rather directly to recent work on indirect or implicit memory. There is also an effortful mode of memory. Effortful suppression of projectable properties allows conceptualization to be guided by trajectories. The resulting conceptualization is what underlies personal, autobiographic, conscious recollection.

5.1 Memory's automatic contribution to conceptualization

When we are walking the path home, we do not need to consciously recall which way to turn at each intersection; when we recognize our children in a crowd, it is not because we have subjected each face to a conscious check; and as we read each word in a sentence, there is no need to try to remember back to when we might have last encountered a similar-looking pattern in order to ascertain the meaning of the word. Memory is automatically, that is, without intention, creating a mesh between the projectable properties (the path, the faces, the letters) and patterns of interaction controlled by non-projectable properties. Research on indirect or implicit memory (Roediger, 1990, 1994; Tulving & Schacter, 1990) is tapping this automatic mode of functioning.

Indirect tests of memory do not require conscious decisions that something is remembered. Instead, the tests often measure some form of repetition priming: The extent to which previous exposure to a stimulus facilitates current processing. For example, a list of words (or pictures) can be presented in phase 1 of a repetition priming experiment. In phase 2, subjects are asked to identify degraded stimuli, some of which occurred in phase 1. Repetition priming is the phenomenon that identification of stimuli actually presented in phase 1 is superior to identification of stimuli presented for the first time in phase 2. This finding occurs whether or not the subjects are attempting to remember anything about phase 1 (e.g., Jacoby & Dallas, 1982; Weldon & Roediger, 1987). A more conceptual form of indirect memory can be measured by, for example, presentation of a word and the later choice of that word as an answer on a test of knowledge (Blaxton, 1989). Among the many interesting findings generated by this research, several may be particularly important. First, repetition priming, can be of very long duration. It is not unusual to be able to demonstrate positive effects over weeks and months (Sloman, Hayman, Ohta, Law, & Tulving, 1988). Second, repetition priming effects are often sensitive to presentation and test modality. For example, pictures prime pictures more than pictures prime words, and vice versa (Weldon & Roediger, 1987). Third, people with dense amnesia often perform equivalent to non-amnesics on indirect, repetition priming tests (e.g., Musen & Squire, 1991). I address this last finding in section 5.2.4.

Jacoby (e.g., Jacoby, Toth, & Yonelinas, 1993) has made the case that much of repetition priming is due to an automatic component. Jacoby characterizes this component as "familiarity" that arises from "perceptual fluency." In the embodiment framework, the automatic component of memory is the contribution of embodied memories to conceptualization of the current environment. It is a type of perceptual fluency in that it affects how aspects of projectable properties are conceptualized. Because embodied memories do not change the clamped environment, the automatic operation of memory does not help one to literally see more clearly (that is, with greater acuity): Instead it helps one to understand the environment. That is why repetition priming has negligible effects on accuracy (in the signal detection sense of ability to discriminate) while affecting interpretation (or bias, Ratcliff & McKoon, 1993; Ratcliff, McKoon, & Verwoerd, 1989).

Repetition priming is modality specific because it is often based on clamped projectable

properties. For example, consider an experiment in which both pictures and words are presented and later subjects must identify the objects in fragmented pictures. To identify the pictured object, subjects must use their memories to mesh with the projectable fragments. Clearly, features of the letters used in spelling the name of the pictured object are irrelevant to this task, so little priming is expected or found between reading words in phase 1 and identifying pictures in phase 2 (Weldon & Roediger, 1987).

Use of trajectories may underlie conceptual forms of repetition priming as well. For example, presentation of "Amazon" in phase 1 will facilitate answering, "What is the longest river," in phase 2. Clearly, words are more than just marks on a page. In reading "Amazon" we think about what rivers are in terms of swimming, fording, etc. This cognitive activity reinforces trajectories from the word "Amazon" to these activities. Later, in comprehending the question, "What is the longest river," we may create a similar conceptualization of rivers. Given the previously reinforced trajectories, the embodied conceptualization of Amazon is easily reachable (that is, meshes with) the embodied conceptualization of "longest river."

5.2 Effortful memory

In Section 3.4 I discussed the idea that multi-step prediction requires suppression to loosen control of projectable properties on conceptualization. I suggested that suppression is dangerous because projectable properties that should be controlling action (such as walking) are ignored. This analysis leads to several suggestions. First, because suppression is dangerous, it is effortful. The effort is a warning signal: Take care; you are not attending to your actions! Also, the effort forces us to use suppression conservatively (because it hurts to think). Second, there are behavioral indices of suppression. For example, when working on a difficult intellectual problem (which should require suppression of the environment), we reduce the rate at which we are walking to avoid injury. Third, autobiographical memory arises from suppressing the environment: Once the environment is suppressed, conceptualization is controlled by trajectories and bodily constraints on mesh rather than the projectable features of the environment. Thus recollection is similar to prediction. Both are effortful, both depend on trajectories, and both are constrained by the body. On this view, conscious recollection is a type of pattern completion that is inherently reconstructive (Bransford, 1979).

The effort in suppressing the environment can be used to explain standard and non-standard facts of episodic memory. As an example of the latter, consider the phenomenon of averting the gaze when engaged in a difficult memory task. When recollection is difficult and unrelated to the current environment, clamping of the environment must be suppressed to allow internal control over conceptualization. Closing one's eyes or looking toward a blank sky are actions that help to suppress the environment by eliminating projectable properties that would normally be clamped. Glenberg, Schroeder, and Robertson (1995) have demonstrated that people avert their gaze when working on moderately difficult recollection tasks (but not easy ones), and that this behavior enhances accurate remembering.

5.2.1 Encoding paradigms

How people are instructed to think about (i.e., encode) to-be-remembered stimuli greatly affects success in conscious recollection. Interactive imagery (e. g., Bower, 1970), levels of processing (Craik & Lockhart, 1972), and generation paradigms (e.g., Slamecka & Graf, 1978) all illustrate this phenomena. As an example, consider the use of interactive imagery to memorize arbitrary pairings such as "lamp - 88." Success in remembering the pairing is greatly enhanced by imagining, say, a neon light shaped to form the digits 88, compared to rote rehearsal of the words.

Standard analyses based on the notion of abstract, amodal symbols have difficulty with these effects because the abstract propositional description of the to-be-remembered stimuli are the same regardless of the encoding task. That is, for both rote rehearsal and imagery one must remember the same thing, "lamp - 88." On an embodied account, constructing an image requires meshing a conceptualization of a lamp with that of 88. The changes in conceptualization from the orthographic stimulus to the meshed image update memory trajectories. Later, partial information such as "lamp" may be given as a cue for the pair. Reading and conceptualizing "lamp" will be along the lines of the reinforced trajectory. Importantly, the analog shapes of the successive conceptualizations increasingly specify the final conceptualization of the neon 88. Contrast this with a situation in which the encoding task is not interactive imagery but simply reading the two words or engaging in rote rehearsal. There is little mesh created by reading the words: the words are pronounced separately so that there is not a physical mesh such as that produced by co-articulation. Furthermore, there is no conceptual mesh in terms of the patterns of interaction between the two objects named by the words. No wonder that little can be reconstructed from the cue "lamp" alone.⁵

5.2.3 The feel of memory

Why is there a phenomenal feel to conscious recollection? Why does the content of memory appear to reflect personal experience? Why doesn't perception or automatic uses of memory feel that way? The feel of memory comes from the effort of suppressing the environment and the consequent knowledge that conceptualization is being driven by previously created trajectories. This process has the feel of personal memory because of our belief that the achieved conceptualization is free from domination by the projectable properties of the environment.

5.2.4 Suppression and amnesia

To the extent that skill in suppressing the environment develops, it suggests explanations for several related phenomena. Consider first infantile amnesia. There is now good evidence that the phenomenon is not as dramatic as initially proposed. In particular, there is evidence for good early retention when it is tested non-verbally. To the extent that a test trades on the automatic operation of memory, it should reveal substantial memory for the infants. In addition, both Howe and Courage (1993) and Nelson (1993) have suggested that what changes around ages 2-3 is the child's ability to code and retrieve information in ways understandable to adults. For Howe and Courage, this amounts to developing a self-concept useful in organizing and retrieving memories. For Nelson, this amounts to learning how to use narrative structures to organize and relate the child's narrative (i.e., self) experiences. Nelson notes that this learning is typically guided by interactions with adults.

⁵ This account is very similar to that given to explain how thinking about "Amazon" can facilitate answering the question "What is the longest river?" That example, however, was used in the context of implicit, or automatic mode of memory. Given that automatic and effortful memory are clearly dissociable, how can the same processes apply? Although dissociable, tasks designed to tap these two modes of memory share many sub-processes (cf. Jacoby et al., 1993). On the embodied account, both modes make use of trajectories to develop a meshed conceptualization: the memory. The purest automatic case is when the trajectories apply to clamped projectable properties, such as fleshing out a fragmented picture. This is what Roediger et al. (1994) call a data-driven task. The purest effortful case is when nothing in the current environment is related to the sought-after information so that all projectable properties must be suppressed. This is an instance of Roediger et al.'s conceptually-driven tasks. The "Amazon" and "lamp" examples are intermediate in terms of suppression. We cannot completely suppress the projectable properties because we need to read the words. Nonetheless, the orthography of the words is of little use; what matters are the trajectories they invoke.

Consider the following explanation for the correlation between development of self-concepts and the emergence of recollective experience. Recollective experience requires a) suppression of environmental input, b) use of self-generated information (trajectories) to drive the conceptual system, and c) an attribution that the resulting conceptualization is due more to internally-guided than externally-guided construction. I suspect that a major factor in the development of a concept of self is just the ability to suppress environmental information. Until that skill is mastered, conceptualization is controlled by the clamped environment; after that skill is mastered, conceptualization can be guided by oneself. That is, one can control what one is thinking about. Furthermore, development of language (by interacting with adults) may well be an important experience in learning how to control suppression and recollective experience: Development of language facility is tantamount to learning to use words to guide conceptualization. Thus skill in suppressing the environment is facilitated by language, and this same skill supports recollective experience and the development of a notion of self.

If recovery from infantile amnesia requires learning to suppress the environment's control of conceptualization, perhaps adult anterograde amnesia results from a traumatically-induced reduction in the ability to suppress. Two findings are consonant with this speculation. First, amnesics exhibit poor performance on explicit tests of memory requiring conscious recollection, but not on implicit tests of memory (e.g., Musen & Squire, 1991). According to the framework developed here, it is the explicit, recollective tests that require suppression of the environment, not the implicit, automatic tests. Second, although there are numerous explanations of amnesic abilities and disabilities, none provides any explanation for the feel of memory. That is, when it can be demonstrated that amnesics are using past experience as effectively as normal rememberers (on implicit or automatic tests), why don't the amnesics have any sense that they are remembering? Of course, the same question can be asked of the normal rememberers: When they perform well on an implicit (automatic) memory task, why do they lack the experience of remembering? For the normal rememberers, the feel of memory comes from an effortful suppression of environmental input and the attribution that conceptualization is controlled by the self. When conceptualization is controlled predominately by the environment, as when performing implicit memory tasks, it does not feel like memory. And, this is the usual state for amnesics.

5.3 The Kolers-Roediger program

Other memory researchers have proposed ideas similar to the framework outlined here. A particularly good example is the "procedures of mind" approach (Kolers & Roediger, 1984). In fact, the similarities between the approaches are striking. Kolers and Roediger suggest that many distinctions popular in memory theorizing reflect different skills rather than different memory stores. Importantly, while championing a symbolic account, Kolers and Roediger note that abstract, meaningless symbols will not do. Instead, they prefer symbols that retain characteristics of how they were acquired: "We claim that knowledge of objects is specific to the means of experiencing them" (page 419). Thus the symbols are in some ways analogical, as I have advocated. Kolers and Roediger also object to modeling knowledge using psychologists' propositions because "descriptions of events rarely if ever tell a person what to do about the events described" (page 439). Of course, conceptualization in terms of patterns of interaction with the environment was designed to overcome this problem. Finally, Kolers and Roediger eschew the idea that memory is purely a conscious experience. They propose instead that the most important contribution of memory is to the automatic execution of skills.

Given the similarities between the Kolers-Roediger program and the embodiment framework, are there any differences? One is my emphasis on meaning, that the meaning of an object or event is a meshed pattern of possible action. A second difference is the idea of mesh

itself. The mesh between the projectable features of an object and non-projectable features from memory can dramatically change the meaning of an object or event (see section 7.3). This sort of combination is made possible by considering both the projectable and the non-projectable features to be patterns of action that can combine as physical, bodily actions can be combined. If separate patterns of action can be forced into a coherent pattern of bodily movement, then we can comprehend the combination; in this way rocks, twigs, and soil combine to form a path for a particular person. It is not clear how the skills described by Kolers and Roediger can be combined except through concatenation. Finally, I am attempting to extend the analysis to language comprehension.

6. Language comprehension

I have argued that the same memory system underlies perception, semantic memory, and episodic memory. The meaning of a situation is given by a meshed pattern of possible actions, and that is an embodied conceptualization. The system is updated whenever there is a change in conceptualization. Thus, the environment is comprehended as a series of transformations of embodied conceptualizations. I propose a similar characterization of language comprehension. Language comprehension, like comprehension of the environment, is the successive transformation of conceptualizations which are patterns of possible action.

Like recollective memory, language comprehension requires suppression of the environment, but in two ways. First, the content of the language may have nothing to do with the physical environment in which the language is expressed. Lectures, for example, have little to do with the lecture hall. Thus, we must suppress projectable properties of the environment. Second, and perhaps more difficult, we must also suppress the projectable properties of the language signal itself. That is, to understand the language, we cannot focus on the shapes of the letters, the patterns of spaces between the words on the page, or the chirps and squeaks of the speech signal.

Several predictions follow from the claim that language comprehension requires suppression of projectable properties. The first is that good language comprehenders should be good at suppressing the environment. Second, good language comprehenders should be good recollectors, given that both require suppression. Third, unavoidable or non-suppressable properties of the environment should disrupt language comprehension. Of course, distracting noise or sights will impair comprehension, but a more subtle effect is discussed by Sanford and Moxey (1995). They note that many types of regularity seem to disrupt language comprehension, and hence those regularities are classified as instances of poor style. Repeating patterns of articulation (McCutchen & Perfetti, 1982), phonemes (e.g., "Crude rude Jude chewed stewed foods," from Baddeley and Hitch, 1974), and excessive repetition of particular sentence structures all seem to slow comprehension. Sanford and Moxey propose that these regularities contribute to the computation of coherence, but because the regularities "are irrelevant to the writer's message," the processing rapidly runs to a halt. Here is a different (but related) suggestion. The regularities are regularities in the projectable properties of the environment. The regularities capture attention and contravene the suppression required for conceptualization and comprehension of what the language is about. That is, instead of paying attention to the meaning of the language, we start to pay attention to the language itself.

Fourth, because language comprehension is seen as a general skill, performance in language comprehension tasks should correlate with performance in other comprehension tasks. Gernsbacher, Varner, and Faust (1990) have demonstrated just this.

6.1 Mental models in language comprehension

Suppose, as Taylor and Tversky (1992) claim, that at least one of the functions of language is that "Language is a surrogate for experience" (page 495). If language is to be a useful surrogate, it must make contact with the sorts of embodied representations that we use to characterize the world, and I propose that language does this relatively directly: We understand language by creating embodied conceptualizations of situations the language is describing. In fact, this is the only reasonable story for how we can manage to learn from language.

This story works when language is being used as a surrogate for events that are completely absent, and when language is being used to enhance current experience. Consider a situation, in which a mother is instructing her child. Representations derived from the language must smoothly integrate (mesh) with representations derived from other aspects of the environment. Thus, being told "That plate is hot," must modify the embodied representation of the plate in order to modify interactions with the plate. Tannenhaus, Spivey-Knowlton, Eberhard, and Sedivy (1995) demonstrated just this sort of smooth and immediate integration. Their subjects responded to verbal commands (e.g., "Put the apple on the towel in the box.") to move actual objects arrayed before them. Eye movements were monitored during the task. Movement of the eyes to referent objects was extremely closely time-linked to the verbal command. Additionally, the environment was used to smoothly disambiguate the language. For example, when considering the language alone, the phrase "on the towel" is temporarily ambiguous. It may describe the location of a particular apple (the apple that is on the towel) or where an apple is to be put. Indeed, when there was only one apple in the array, the eye movements indicated uncertainty. When the array contained two apples (one on a towel and one on a napkin), however, then the phrase "on the towel" will almost certainly be meant to specify a particular apple, not a location in which to put the apple. In the two-apple case, the eye movements indicated no uncertainty. Thus, understanding of the sentence made virtually immediate use of the context, in contrast to notions of modularity of syntactic analysis. This sort of integration is possible if both the environment and the language are understood as embodied patterns of action.

This sort of reasoning is compatible with work on mental model theory. The basic claim of mental model theorists is that language comprehension results in representations of what the language is about, not representations of the language itself (e.g., chirps, words, sentences, or propositions). Johnson-Laird (1989, p. 488) writes that a mental model is a representation of a situation so that "its structure corresponds to the structure of the situation that it represents." With an important emendation, this definition can apply to the sorts of representations I have been describing. It seems unlikely that the literal, in-the-head structure of the representation could actually be isomorphic to the structure of the situation (in contrast to Glenberg, Kruley, & Langston, 1994).

6.1.1 Mental models from language and perception

Embodied mental models are "models" in the following sense: A model is useful if it can be used to predict the effect of an action in the real situation being modeled. One way to ensure accurate prediction is to build into the model spatial-functional constraints analogous to those of the real situation. For example, a useful model of an airplane will have wings that generate lift when it interacts with air currents, much like a real airplane's wings generate lift. Similarly, a mental model built from language incorporates embodied constraints on action like those derived from comprehension of the environment. This sort of mental model is useful to the extent that it incorporates enough constraints on action to derive predictions.

One difference between embodied models derived from language and those derived from perception is how completely the meshed pattern of possible actions constrains further action

and prediction. The multiple projectable properties of the environment, because they are clamped, tightly constrain conceptualization and action. In language comprehension, the patterns of possible action that contribute to a meshed conceptualization are much looser. That is, language is ambiguous in not specifying exact parameters of spatial layout, force, etc. (Talmy, 1988). Thus, conceptualizations derived from language do not constrain action as effectively as conceptualization derived from the environment. This is one reason for differences between expert and non-expert comprehension. The expert's model incorporates tighter constraints on action based on trajectories derived from experience. Given the same text, the expert is able to take (appropriately constrained) actions that leave the non-expert baffled. This effect of expertise in language comprehension parallels the expert guide who can spot the trail (based on trajectories derived from experience) while the novice sees only twigs, soil, and rocks.

6.2 Comprehension, prediction, and priming

I have argued that embodiment in terms of action patterns is just what is needed to facilitate interaction with the environment and prediction. Is prediction an important component of language comprehension? Clearly, language would be of little use if it did not enable better prediction of the environment. But, the question asked in the literature on comprehension is different: Does a mental model serve as a source of "on-line" predictions about the upcoming text? In response to this question, one might ask "Should it?" If the point of language is to be a surrogate for experience, that is, to help us take appropriate action in real situations, it makes little sense to expect the representation to predict upcoming text: It should predict changes in the situation. In fact, McKoon and Ratcliff (1992) reached the conclusion that there is little evidence that people make predictive inferences while reading.

Nonetheless, Keefe and McDaniel (1993) presented convincing evidence for what appeared to be just those inferences. Using the standard logic based on psychologists' propositions, Keefe and McDaniel reasoned that pronunciation of a probe word would be faster if the word were part of a recently made inference than if not. For example, subjects read a sentence such as, "After standing through the three-hour debate, the tired speaker walked over to his chair." Following the sentence, subjects pronounced the probe word "sat." Supposedly, pronunciation of the probe word would be facilitated by its having been incorporated into an inferred proposition such as "The speaker sat down." In the control condition, for which an inference including the word "sat" is unlikely, subjects read a sentence such as, "The tired speaker moved the chair that was in his way and walked to the podium to continue his three-hour debate." Indeed, pronunciation of the probe word was faster in the predictive condition than in the control condition. In fact, pronunciation of the probe following the predictive sentence was as fast as when the sentence explicitly continued with "and (he) sat down." Murray, Klin, and Meyers (1993) used a similar methodology and produced a similar effect when the "to-be-inferred event was in focus at the time of test" (page 464). Why is evidence for predictive inferences found only shortly after the predicting sentence? Does this evidence demonstrate that subjects were attempting to predict the upcoming text?

Consider an interpretation of these findings from the point of view that the goal of language comprehension is the creation of a conceptualization of meshed patterns of action. In this case, interpretation of a word, phrase, or sentence, consists of meshing the actions consistent with that bit of language with the patterns of action derived from previous text. After comprehending Keefe and McDaniel's predictive sentence ("...the tired speaker walked over to his chair"), only certain actions can be easily meshed with the conceptualization. For example, the actions implied by "He began to rake the leaves," does not mesh. In contrast, the action of sitting will mesh, and hence interpretation and pronunciation of the probe word "sat" is quick. After comprehending Keefe and McDaniel's control sentence ("...walked to the podium to

continue his three-hour debate") the action of sitting meshes about as well as the action of raking, and so pronunciation of the probe word "sat" is slow.

This interpretation of the results is radically different from that used in standard propositional accounts of inference making. In the standard account, an inference corresponds to encoding a new proposition, something akin to, "He sat down," and one would expect some effect of this proposition well after it was encoded. The embodied account is that no "inference" in the standard sense is made. Instead, the action of sitting in the chair is temporarily compatible with the embodied conceptualization. When the situation changes, some actions are no longer compatible with the embodied conceptualization and the "inference" is no longer operative. This notion of temporary compatibility (how well the probe will mesh with the other constraints) may well underlie McKoon and Ratcliff's (1986b) data for "partial" encoding of predictive inferences, and the temporary effect noted by Keefe and McDaniel (1993) and Murray et al. (1993). Of course, this is not to say that language comprehenders might not make forward inferences if induced to do so (e.g., one might be asked to "Guess what happened next"). These sorts of inferences are just the sort of predictions discussed in Section 3.4. However, given that language, unlike the environment, only loosely constrains action, it is more reasonable to wait until what happens next is described.

The procedures used by Keefe and McDaniel (1993) and Murray et al. (1993) follow from the more general notion of "semantic priming." The standard idea is that processing causes activation to spread along permanent links to semantically related information, and this spread of activation speeds processing of the related information. Thus, reading the prime, "doctor" speeds the decision that the target "nurse" is a word (Meyer & Schvaneveldt, 1971). The standard interpretation of semantic priming is embarrassed by demonstrations that priming need not be due to permanent links (McKoon & Ratcliff, 1986a), and that the effective relation between the prime and the target may have little to do with the presumed semantics of categories (Sheldon & Martin, 1992). Might semantic priming be another instance of the operation of mesh? Assume that language comprehension is an attempt to mesh action suggested by the current word or phrase with the pattern of actions already established. Thus, in thinking about what a "doctor" is (the actions taken by a doctor and how one interacts with a doctor), one sets up a conceptualization in which the actions suggested by "nurse" will mesh. Hence, processing of "nurse" is facilitated relative to the case when it is preceded by an unrelated prime word such as "rake."

A report by Hess, Foss, and Carroll (1995) strongly suggests that semantic priming reflects something akin to mesh rather than spread of activation along permanent links. Their subjects read a sentence describing a local context, such as "To complete the assignment, the English major wrote a..." and they then read a target word such as "poem." The question of interest was whether the local context ("English major") would facilitate reading of the target ("poem") regardless of the global situation. This would be expected if priming reflects activation along permanent links such as between "English major" and "poem." In one global situation, the English major was working on a writing assignment, and indeed, reading of "poem" was facilitated relative to a neutral condition. In another global situation, however, the English major was working on a computer program. In this case, reading the target "poem" was not facilitated. The implication is that priming reflects ease of integration (mesh) of concepts, not spread of activation along permanent links.

6.3 Space in language comprehension

If embodied conceptualization is a pattern of possible actions, then it must incorporate information about spatial layout, because actions are played out in space. The data from several research projects investigating spatial coding in mental models provide this evidence. First,

there have been investigations of how language can lead to accurate, analogical representations of a described layout. For example, Denis and Cocude (1989) had subjects read texts describing the layout of objects on a circular island. After several readings, they were asked to mentally simulate scanning from one object to another. The main finding was of a correlation between distance (if the objects had actually been arrayed) and simulated scanning time. Morrow, Bower, and Greenspan (1989, see also Rinck & Bower, 1995) had their subjects memorize the layout of the rooms in a building (and objects in the rooms) before reading a passage describing the movements of a protagonist throughout the building. Morrow et al. measured time to verify that particular objects were located in particular rooms as a function of the protagonist's described movements. Interestingly, when a described movement (e.g., from Room A to Room C) required passage through an unnamed room on the path of the movement, verification of objects located in the unnamed room was faster than verification of objects in other unnamed rooms off the path. Apparently, subjects were using the spatial information in the building layout while comprehending the text.

Glenberg, Meyer, and Lindem (1987) demonstrated the contribution of spatial information to language comprehension without pre-memorization. Subjects read texts describing a protagonist and a target (e.g., a jogger and a sweatshirt) that were either spatially dissociated (the jogger took off his sweatshirt before jogging) or spatially associated (the jogger put on his sweatshirt before jogging). After a sentence or two in which the protagonist was kept foregrounded but the target was not mentioned, accessibility of the target (e.g., the sweatshirt) was greater in the associated condition than in the dissociated condition. [McKoon and Ratcliff, 1992, have argued that this effect may reflect a type of salience. See Glenberg and Mathew (1992) for a counter to this interpretation.] Along similar lines, O'Brien and Albrecht (1992) demonstrated sensitivity to spatial location of characters in a text well after the spatial information is introduced. Thus, several sentences after reading, "As Kim stood outside the health club, she felt a little sluggish," readers would balk at the sentence, "She decided to go outside..."

One interpretation of these findings is that they reflect a representation that is analogical with respect to space, that is, that the mental model is constructed in an inherently spatial medium. This seems unlikely. Langston, Kramer, and Glenberg (1995) have demonstrated that spatial contiguity, in the absence of other relations, does not have strong functional consequences. In these experiments, subjects read (or heard) texts describing the spatial layout of four objects. In outline, the texts read, "B is to the right of A, C is under B, D is to the left of/right of C." The last sentence in the "close" condition was "D is to the left of C," so that the spatial layout of the objects has D under (that is, close to) A. The last sentence in the "far" condition was "D is to the right of C," so that object D is separated from A. After reading, subjects were tested for availability of the target object, A. If space is represented analogically, and if closeness in that space has functional consequences, then the target should be more available in the "close" condition than in the "far" condition. We tested for availability of A using speeded recognition of A and time to read a sentence referring to object A. Availability of A was never reliably affected by the condition ("close" versus "far"), even though memory for the spatial layout was well above chance.

How are we to understand the contrast between Langston et al. (1995) and the other research that clearly points to an appreciation of spatial relations during comprehension? One possibility builds on the distinction between mental models encoding space in a spatial medium and mental models encoding spatial-functional action and thereby representing space incidentally. Consider a reinterpretation of Glenberg et al.'s (1987) jogger on this spatial-functional account. When the jogger puts on the sweatshirt, there is a mesh between the jogger and the sweatshirt: Wherever the jogger goes, the sweatshirt goes too. Then, later facilitation in reading "sweatshirt" is not due to spatial closeness of the jogger and the sweatshirt, but their

functional relatedness. On this account, the texts used by Langston et al. (1995) resulted in encoding patterns of action between the reader (projected into the situation) and each object (A, B, C, and D). Given that spatial layout is not encoded directly, there is little reason to suspect that availability of object A will depend on its spatial distance from object D. In other words, spatial distance only matters when it corresponds to functional distance.

The proposal that embodied mental models reflect a structured space (that is, a space structured by possible actions) rather than a uniform space, is consistent with several research programs. McNamara (1986) and McNamara, Hardy, and Hirtle (1989) adduce evidence that spatial memory is structured and perhaps hierarchical. Bryant, Tversky, and Franklin (1992) argue that the time needed to answer questions about memorized spatial layouts reflects an embodied encoding. They find that retrieval of information aligned on the head/feet axis is faster than for the front/back axis which is faster than for the left/right axis. They interpret these differences as reflecting asymmetries of the body.

6.4 Comprehension of non-concrete descriptions

If language comprehension is in terms of meshed action, how is it that we come to understand abstract language that is not about concrete objects or situations? Here I adopt a version of Lakoff's (1987) spatialization of form hypothesis. Namely, we understand abstract situations by conceptualizing them in concrete ways.

Talmy's (1988) analysis of force dynamics is a good example of how abstract concepts can be given a bodily interpretation. He notes that we can conceptualize forces as one entity (an agonist) acting against another (an antagonist) and that the entities may have different strengths and different tendencies (either toward action or toward inaction). Importantly, these basic entities and relations can be based on bodily experiences such as pushing and being pushed, moving objects, etc. Talmy suggests that our understanding of causal terms (e.g., "because") reflects an agonist's tendency (toward action or inaction) being overcome by an antagonist. Thus, we understand the sentence "The ball kept rolling because of the wind blowing on it" as an agonist (the ball) with a tendency toward inaction being overcome by the stronger antagonist, the wind. Talmy also demonstrates how this analysis can be extended to psychological instances of causation, social references, and interpretation of modals such as can, may, must, and should. Thus, interpretation of "John cannot leave the house" comes about from assigning John the role of an agonist whose actions are blocked by the unmentioned but stronger antagonist of social or physical constraint. In the case of "should not" the antagonist is a value or belief, and so on. The point is that what has traditionally been treated as prototypically abstract (e.g., cause, force, modality), can be conceptualized in embodied terms, and in so doing brings out important similarities in our understanding of these concepts.

Bowerman (1982, 1985) discusses a number of cases of children's late speech errors that imply an understanding of the more abstract in terms of the concrete. Bowerman classifies an error as a late speech error when it occurs after a linguistic form has been used correctly and when the error does not mirror adult usage. She argues that given these constraints, the error arises from an over-extension of the adult-sanctioned relation between domains. Typically, the spatial domain is extended, so that the pattern of errors is asymmetrical across domains. For example, children import spatial terms into other domains, but rarely vice versa. Bowerman reports that children use the spatial verbs "put" and "take" to describe state changes, such as "put the door locked." Also, it is commonplace to use spatial terms when describing time, e.g., "the week before," and "between spring break and finals week." Is this just a convention, or does it reflect a conceptualization in which we understand time by using spatial dimensions? The late error, "behind the dinner" for "after the dinner" would seem to imply the latter (Bowerman, 1982, 1985). Finally, Bowerman describes the use of spatial terms to speak of non-spatial

dimensions, such as looseness of teeth ("They're all the same length of loose") and temperature of water ("I want it the same size as Christy's was").

One final example should suffice. Suppose that we conceptualize abstract trait information (e.g., that Marta is energetic), as a meshing between the person and the trait. That is, the actions that Marta might perform are meshed with "energetic" so that her actions are constrained to be energetic. To test this notion, Fernandez and Saiz (1989) had subjects read texts describing the association or dissociation of a main character and a trait. In a text about Marta, an expert in international business, the critical sentences in the associated condition read (in translation from the Spanish):

(1) She has just been appointed to a government position. Almost everybody considers her an especially energetic person.

Whereas the critical sentences in the dissociated condition read:

(2) She has just been appointed to a government position. Almost nobody considers her an especially energetic person.

After reading one or two filler sentences in which Marta was kept foregrounded (but her energy never mentioned), accessibility of "energetic" was evaluated by speeded recognition of the probe "energetic." On average, responding in the associated condition was over 100 msec faster than responding in the dissociated condition. Thus, readers may well have been conceptualizing abstract trait information as embodied and meshed with an embodied conception of Marta.

6.5 Embodiment and coherence

Some texts make sense; others do not. The ones that make sense are judged coherent. But, what produces that sense of coherence? A standard answer is that it arises from the connectedness of the psychologists' propositions underlying the text; when the propositions are connected (or can be made connected through bridging inferences, Haviland & Clark, 1974) then the text is coherent. When the propositions do not connect, either bridging inferences need to be made to connect them, or the text will appear incoherent.

This interpretation of coherence is wrong in several respects (see Sanford and Moxey, 1995). Importantly, the account is wrong because whether or not propositions connect and how they connect depends first on interpreting the propositions against a situation. Consider the following example adapted from Sanford and Moxey:

(3) While measuring the wall, Fred laid the sheet of wallpaper on the table. Then he put his mug of coffee on the wallpaper.

(4) After measuring the wall, Fred pasted the wallpaper on the wall. Then he put his mug of coffee on the wallpaper.

A propositional analysis does not reveal that (4) is odd, and thus a propositional analysis cannot indicate "local incoherence" and cannot trigger bridging inferences to maintain coherence (see also O'Brien & Albrecht, 1992). Noticing that (4) is odd arises from a consideration of the situation, that once the wallpaper is on the wall, under normal conditions, it cannot not support a mug of coffee. To state it differently, coherence is a relationship among ideas, and texts do not have ideas, only readers do.

An impressive counter to the claim that coherence derives from connecting propositions can be found in Barton and Sanford (1993). Their subjects read about an airplane crash that occurred in the Pyrenees between France and Spain. The subjects were asked for advice on where the survivors should be buried. In fact, the subjects readily offered advice; that is, they understood the text, judged it as coherent, and were ready to suggest where the survivors should be buried. Nonetheless, only about 60% of the readers noticed that "survivors" are not buried. In a second experiment, when readers were asked where to bury the "surviving dead," only 23% noticed a problem. Clearly, the readers were not forming propositions and checking them for sensibility, because "surviving dead" cannot make a sensible proposition.

An alternative account of coherence is twofold. First, coherence is a matter of degree, and in fact, no bit of language is completely incoherent. Second, the degree of coherence can only be computed from the mesh of a situational representation of what the language is about.

The claim that no bit of language is completely incoherent rests on the analogy between understanding language and understanding the environment. Consider, for example, a percept of a drawing of an "impossible object." There may be no three-dimensional object that could project that two-dimensional outline. Nonetheless, the percept is not incoherent; the percept is of a drawing that has no corresponding three-dimensional realization. Percepts may be unusual or bizarre, but never incoherent because the perceptual/action system is designed to transduce patterns of possible interaction. Similarly, a random collection of words (or even phonemes or features) will be perceived coherently, perhaps as chirps and whistles, and a random collection of sentences will be perceived coherently (correctly) as a random collection of sentences.

Nonetheless, we do get the sense that some collections of sentences are not random. Sentences cohere to the extent that they produce continuous transformations (trajectories) of a meshed set of possible actions. Consider (4) again. The second sentence seems incoherent in that it cannot be incorporated into the the standard situational interpretation of flat wallpaper on a vertical wall in a gravitational field. However, if the initial model is changed so that any of these presuppositions about the situation are eliminated (e.g., the wallpaper has niches in it, the wall is not yet vertical because it will be incorporated into a doll's house, etc.) then the sentences are coherent. Another example is also adapted from Sanford and Moxey:

(5) John ate a banana. The banana was brown. Brown is a good color for hair. The hair of a dog is drink to counteract a hangover.

Sanford and Moxey use this snippet of text to illustrate that sentences that incorporate cohesion markers (e.g., anaphoric reference) can, nonetheless, be judged incoherent. The problem is that the sentences do not update a mental model. That is, the patterns of action suggested by each sentence do not admit to smooth transformation of the mesh from one sentence to the next. Note, however, that as with the previous example, a change in the initial situation can render the sentences (more) coherent. Imagine that John engages in free association whenever he eats fruits. Then, the list of sentences, as descriptions of his free-associations, seem (more) coherent. Similar examples can be constructed for film (e.g., the sequence of cuts seem incoherent unless one has the appropriate model of the film) and for events in the world (e.g., changes in the weather seem incoherent unless one has the appropriate model of weather systems). In short, coherence is a property of models (the ideas that people have), not a property of snippets of language.

As a final example, (6) was taken from the abstract of a talk given in a Computer Sciences seminar at the University of Wisconsin-Madison.

(6) The talk will concentrate on the design of the communications subsystem [of the

Meiko CS-2 MPP System]. This utilizes a 'fat tree' network constructed from high performance crosspoint switches. Processing Elements interface to this network via a communications co-processor which contains intelligence to handle virtual addressing and ensures very low message start up times.

This text may well be very coherent for its intended audience, but it is at the low-end of the dimension for me. The problem is not that the propositions do not connect. The propositional relation between 'fat tree' networks and crosspoint switches is virtually transparent; similarly, it is quite clear that a co-processor intervenes between the "Processing Elements" and the network. The problem is that I do not know what a 'fat tree' network is, or what "crosspoint switches" or "Processing Elements" are. I do not know the literal shapes of these things, nor do I know the actions they can take or how I can interact with them. Because I lack that knowledge, I cannot build a coherent spatial-functional model. Presumably, crosspoint switches can be arrayed or interconnected in some way so that they comprise a 'fat tree' network. But for me, the mesh is missing.

The ideas a) that coherence is a function of the mesh in an embodied model, b) that the embodied models constructed to understand language are the same as those that underlie comprehension of the natural environment, and c) that the purpose of perception and memory for the natural environment is to guide action, all lead to a suggestion about how to assess comprehension. Most laboratory comprehension tests require verbatim reproduction of a text, reproduction of "idea units," or speeded responding to words or phrases. A more sensible comprehension test, however, is one that requires action. To what extent can the reader take sensible action (or make sensible predictions) on the basis of the text? (6) is relatively incoherent for me because I can make so few predictions. For example, if the type of switches were changed, I don't know if that would change the network from a 'fat tree' network to some other kind; if the communications co-processor was not intelligent, I do not know if the message start up times would be slower or faster. On the other hand, (6) is not completely incoherent because there are some predictions that I can make. For example, based on knowledge of part-whole relations, I can predict that if the crosspoint switches are eliminated, there will be no 'fat tree' network.

7. Conclusions

I began with a consideration of the Lakoff and Johnson program and the problem of meaning. In applying their insights to a theory of memory and mental models, the concept of embodiment becomes central. The basic claim is that an individual's memory serves perception and action. Memory meshes non-projectable features with projectable features of the environment to suggest actions for that person in that situation. These patterns of action are what make the environment meaningful to that person. This framework provides a way to address meaning, symbol grounding, recollective and automatic uses of memory, and language comprehension.

7.1 Summary of interpretations and predictions

The framework provides alternative accounts of standard phenomena and it makes new predictions. Here is a brief review. The concept of embodied knowledge is used to address the problem of meaning and symbol grounding (Section 1.3), why people see the world differently (1.3), effects of bodily activity on emotions (2.3.1), imagery (2.3.2), memory for actions (2.3.3), sensibility judgments (2.3.3), short-term behavior (4.3), and understanding in abstract domains (6.4). Mesh of patterns of action is applied to emergent features of thought (3.1), recollective memory (5.2), interactive imagery (5.2.1), interpretation of semantic priming phenomena including forward inferencing (6.2), and coherence (6.3). Suppression of

projectable properties of the environment is seen as critical to multi-step prediction (3.4), the feeling of memory (4.0), the decrease in physical activity when thinking (5.2), amnesia (5.2.4), correlation of language comprehension with recollection (6.0), and effects of incidental patterns on comprehension (6.0). Finally, trajectories are applied to frequency effects in memory (3.3), the nature of rehearsal (4.3), automatic uses of memory (5.1), and expertise (6.1.1).

7.2 Embodied knowledge, emotions, and social behavior

Can embodied patterns of action underlie all conceptualization? Our experiences of music, taste, and emotions all seem to have aspects that do not fit well into a spatial-functional straitjacket, and one suspects that aspects of these experiences are represented in addition to action patterns. Nonetheless, given the ease with which these sorts of experience combine with spatial-functional experience (consider the contribution of music and mood to the understanding of the action depicted in a film), it is not inconceivable that they may eventually be covered by the same sort of analysis.

Missing from the discussion is a consideration of hedonic valence and motivation to act. It is not as yet clear how pleasure and pain should be represented in an action-oriented system (but see Lang, 1979). What is clear, however, is that hedonic valence affects action and how experiences become meaningful. Our understanding of pleasurable experiences is in part action-toward those experiences, whereas our understanding of aversive experiences is in part action-away. Several ideas follow. Given that action-away does not necessarily specify what the action is directed toward, it ought to be more diffused and variable than action toward. Also, on this analysis, approach and withdrawal are not poles of a single dimension: Withdrawal from one situation does not imply approach toward another. Thus, our understanding of emotional experience should reflect at least two dimensions (e.g., Schneirla, 1959).

Malter (1996) applies these ideas to consumer research, in particular, impulse buying. He proposes that projectable features of a product automatically mesh with affectively charged memories (perhaps imparted by advertisements) to produce an irresistible approach-dominated conceptualization. Thus the consumer experiences a strong desire to approach and manipulate the object, and in most cases that can only be accomplished after purchase. Furthermore, Malter notes that overcoming this urge to buy requires effortful suppression of the projectable features in order to deliberately evaluate the purchase. In the face of a strong impulse to buy, however, that effort may be viewed as unattractive or not considered at all.

There is also reason to believe that an embodied, action-oriented analysis has implications for social psychology. Fiske (1992) traces the history of action-oriented theories of social cognition from James ("My thinking is first and last and always for the sake of my doing..." as quoted in Fiske) to Heider (1958) to current "pragmatic" research. Fiske defines pragmatism as a framework in which "meaning, truth, and validity are determined by practical consequences [and] concrete goal-relevant actions," (p. 886).

Fiske's own analysis of social cognition is compatible with the ideas I have described, and her analysis suggests an important extension. According to Fiske (and Heider, 1958), the key to social cognition is to view others not just as objects that we can affect, but as beings who can effect us in turn. Consonant with this premise, Fiske proposes that our ability to infer traits is in the service of interaction with others.

7.3 Mesh

Among the interpretations that stem from a consideration of embodied representations, the notion of mesh seems most important. Ideas mesh to the extent that the pattern of action

underlying one idea can be integrated with the pattern of action underlying another. The patterns mutually modify and constrain one another because the conjoint actions must be possible given our bodies. This mutual modification of patterns of action is what underlies the construction of meaning from words whose senses are jointly modified by the contexts in which they occur.

Meshing patterns of action provides a new way of thinking about componentiality and productivity in language. As an example, consider the coke bottle. Its shape, and thus its affordances for human action, allow it to mesh with many physical situations and goals. It can be used for storing liquid, as a cup, a doorstop, a weapon, a vase, and so on. Thus the meaning of a coke bottle (how we can interact with it) is not fixed, but infinitely varied, depending on the context of use. Importantly, however, the meaning is in no way arbitrary or unconstrained: The meaning of the bottle is constrained by its shape (heft, fragility, etc.) and the implications of that shape for action. Thus the spatial-functional meaning of a coke bottle is componential in that it will mesh with many human contexts. Because that mesh can transform the meaning, however, its use is creatively productive.

This type of componentiality helps us to understand what Barsalou et al. (1993) term "linguistic vagary." When people are asked to describe the features of a category such "coke bottle" there is tremendous variability both across people and from time to time in a particular person's descriptions. Linguistic vagary should be the norm if the meaning of a concept is determined by its mesh with the context.

The idea of mesh may prove to be a concept that can replace "association." Although association has played a central role in theories of cognition, the term carries little theoretical weight. What we mean by an association is little more than a conditional probability; if B is associated with A, then $P(B|A) > P(B)$. There is little or nothing in our theories to help us understand when "laws of association," such as frequency and recency, hold, and when they do not. In contrast, the notion of mesh can, like an association, be used to relate concepts, but the nature of the relation is deeper: When patterns mesh, they modify each other because they must conjoin in a way that respects constraints on bodily action. Thus, "coke bottle" is difficult to mesh with "chair."

Mesh provides a rationale for Thorndike's (1932) concept of belongingness (see also, Ohman, Fredrikson, Hugdahl, & Rimmo, 1976) as well as various ideas put forward by Gestalt psychologists. Furthermore, mesh may help to explicate species-specific differences in associability of stimuli. Rats find it easier to associate a novel taste with illness than to associate a novel sight with illness (Garcia & Koelling, 1966). In contrast, pigeons find it easier to associate a novel sight, rather than a novel taste, with illness (Wilcoxin, Dragoin, & Kral, 1971). If learning comes about through meshing patterns of bodily action, then, given species differences in anatomy, physiology, and possible actions, the fact that stimuli will mesh differently for different species is a foregone conclusion.

7.4 Standard memory paradigms

If knowledge is embodied, then commonly-used laboratory paradigms for studying memory may well be missing the mark. Many of these paradigms use random lists of words as the objects of memory. Whereas there are reasons for using lists of words, these are reasons related to history and convenience, not to any analysis of the design of memory. A favorite argument to justify the verbal list format is that each word corresponds to a mini-event, and from memory's point of view these mini-events are similar to other events in the world. This argument loses much of its force, however, if memory is embodied and designed for negotiating a three-dimensional environment.

The discrepancy between the design of memory and the design of the tools used to analyze it may account for undesirable characteristics of memory research. Importantly, memory researchers have not made much progress in understanding the nature of memory. We know about many phenomena (Greene, 1992), but there is little agreement as to the interpretation of those phenomena, how they fit together, or whether a particular phenomenon is of any importance. Even with something as basic as the effect of repetition, the theoretical diversity is astounding: We have theories in which repetitions enhance the strength of a single representation (Gillund and Shiffrin, 1982), theories in which repetitions are individually preserved (Hintzman, 1986), and theories which treat memory much like a hologram (Metcalf, 1993). We have multi-store theories and single-store theories; single system theories and multiple-system theories. All of these positions receive support from some aspects of the literature. I suspect that this diversity of positions arises because in using inappropriate tools we obtain incompatible views of memory much like the views of the blind men touching the elephant.

References

- Baddeley, A. (1990). *Human memory*. Boston: Allyn and Bacon.
- Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In G. Bower (Ed.), *The psychology of learning and motivation*, Vol. 8 (pp. 47-89). New York: Academic Press.
- Barsalou, L. W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A. C. Collins, S. E. Gathercole, & M. A. Conway (Eds.), *Theories of memories* (29-101). London: Erlbaum.
- Barsalou, L. W., Yeh, W., Luka, B.J., Olseth, K. L., Mix, K. S., & Wu, L-L. (1993). Concepts and meaning. In K. Beals, G. Cooke, D., Kathman, K. E. McCullough, S. Kita, & D. Testen (Eds.) *Chicago Linguistics Society 29: Papers from the parasession on conceptual representations*. University of Chicago: Chicago Linguistics Society.
- Barton, S. B., & Sanford, A. J. (1993). A case study of anomaly detection: Shallow semantic processing and cohesion establishment. *Memory & Cognition*, 21, 477-487.
- Berkowitz, L., Jo, E., & Troccoli, B. (1993). Feeling and knowing: Evidence for different psychological processes in emotional reactions. Manuscript in preparation.
- Berkowitz, L. & Troccoli, B. T. (1990). Feelings, direction of attention, and expressed evaluations of others. *Cognition and Emotion*, 4, 305-325.
- Blaxton, T. A. (1989). Investigating dissociations among memory measures: Support for a transfer-appropriate processing framework. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 657-668.
- Bloom, L. (1994). Meaning and Expression. In W. F. Overton and D. S. Palermo (Eds.), *The nature and ontogenesis of meaning* (pp.215-235). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Bower, G. H. (1970). Analysis of a mnemonic device. *American Scientist*, 58.
- Bowerman, M. (1982). Starting to talk worse: Clues to language acquisition from children's late speech errors. In S. Strauss (Ed.), *U-shaped behavioral growth* (pp. 101-145). New York: Academic Press
- Bowerman, M. (1985). *The spatial metaphor: A universal principle of semantic structuring? Clues from children's speech errors*. Oxford University Psychology Department Colloquium.
- Bransford, J. D. (1979). *Human Cognition: Learning, Understanding and Remembering*. Belmont, CA: Wadsworth Publishing Company.
- Bruce, V., & Green, P. R. (1985). *Visual perception : physiology, psychology and ecology*. Hillsdale, N. J.: Erlbaum.
- Bryant, D. J., Tversky, B., & Franklin, N. (1992). Internal and external spatial frameworks for representing described scenes. *Journal of Memory and Language*, 31, 74-98.
- Carpenter, P. A., Miyake, A., & Just, M. A. (1994). Working memory constraints in comprehension: Evidence from individual differences, aphasia, and aging. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*. Orlando, FL: Academic Press.
- Cohen, R. L. (1981). On the generality of some memory laws. *Scandinavian Journal of Psychology*, 22, 267-281.
- Collins, A. M., & Quillian, M. R. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8, 240-247.
- Conrad, C. (1972). Cognitive economy in semantic memory. *Journal of Experimental Psychology*, 92, 149-154.
- Craik, F. I. M. & Lockhart, R. S. (1972). Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior*, 11, 671-684.
- Cuthbert, B. N., Vrana, S. R., & Bradley, M. M. (1991). Imagery: function and physiology. In *Advances in psychophysiology*.
- Denis, M., & Cocude, M. (1989). Scanning visual images generated from verbal descriptions. *European Journal of Cognitive Psychology*, 1, 293-307.

- Engelkamp, J., & Krumnacker, H. (1980). Imaginale und motorische Prozesse beim Behalten verbalen Materials. *Zeitschrift für experimentelle und angewandte Psychologie*, 27, 511-533.
- Epstein, W. (1993). The representational framework in perceptual theory. *Perception & Psychophysics*, 53, 704-709.
- Fernandez, A., & Saiz, C. (1989, July). *Foregrounding of trait information during text comprehension*. Paper presented at the First European congress of Psychology, Amsterdam, NL.
- Fiske, S. T. (1992). Thinking is for doing: Portraits of social cognition from daguerrotype to laserphoto. *Journal of Personality and Social Psychology*, 1992, 63, 877-889.
- Franklin, N., Tversky, B., & Coon, V. (1992). Switching points of view in spatial mental models. *Memory & Cognition*, 20, 507-518.
- Friedman, W. J. (1993). Memory for the time of post events. *Psychological Bulletin*, 113, 44-66.
- Garcia, J., & Koelling, R. A. (1966). Relation of cue to consequence in avoidance learning. *Psychonomic Science*, 4, 123-124.
- Gärdenfors, P. (1995). *Cued and detached representations in animal cognition*. Manuscript submitted for publication.
- Gernsbacher, M. A., Varner, K. R., & Faust, M. E. (1990). Investigating general comprehension skill. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16, 430-445.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. New York: Houghton Mifflin.
- Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, 91, 1-67.
- Glenberg, A. M. (1984). A retrieval account of the long-term modality effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 16-31.
- Glenberg, A. M., Kruley, P., & Langston, W. E. (1994). Analogical processes in comprehension: Simulation of a mental model. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics*. Orlando, FL: Academic Press.
- Glenberg, A. M., & Mathew, S. (1992). When minimalism is not enough: Mental models in reading comprehension. PSYCOLOQUY.92.3.64.reading-inference-2.1.
- Glenberg, A. M., Meyer, M., & Lindem, K. (1987). Mental models contribute to foregrounding during text comprehension. *Journal of Memory and Learning*, 26, 69-83.
- Glenberg, A. M., Schroeder, J., and Robertson, D. (1995). Remembering behavior. Manuscript in preparation.
- Glenberg, A. M., Smith, S. M., & Green, C. (1977). Type I rehearsal: Maintenance and more. *Journal of Verbal Learning and Verbal Behavior*, 16, 339-352.
- Glenberg, A. M., & Swanson, N. C. (1986). A temporal distinctiveness theory of recency and modality effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 3-15.
- Greene, R. L. (1986). Sources of recency effects in free recall. *Psychological Bulletin*, 99, 221-228.
- Greene, R. L. (1992). *Human memory*. Hillsdale, NJ: Lawrence Erlbaum and Associates.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42, 335-346.
- Harnad, S. (1993) Grounding Symbols in the Analog World with Neural Nets. Think (Special Issue on Machine Learning)
- Haviland, S. E., & Clark, H. H. (1974). What's new: Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior*, 1974, 13, 512-521.
- Heider, F. (1958). Perceiving the other person. In R. Tagiuri & L. Petrullo (Eds.), *Person perception and interpersonal behavior*. Palo Alto, CA: Stanford University Press.
- Hess, D. J., Foss, D. J. & Carroll, P. (1995) Effects of global and local context on lexical

- processing during language comprehension. *Journal of Experimental Psychology: General*, 124, 62-82.
- Hintzman, D. L. (1967). Articulatory coding in short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 6, 312-316.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, 93, 411-428.
- Howe, M. L., & Courage, M. L. (1993). On resolving the enigma of infantile amnesia. *Psychological Bulletin*, 113, 305-326.
- Hunt, R. R., & Elliott, J. M. (1980). The role of nonsemantic information in memory: Orthographic distinctiveness effects upon retention. *Journal of Experimental Psychology: General*, 109, 49-74.
- Jacoby, L. L., & Dallas, M. (1981). On the relationship between autobiographical memory and perceptual learning. *Journal of Experimental Psychology: General*, 110, 306-340.
- Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, 122, 139-154.
- Johnson, M. (1987). *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago: The University of Chicago Press.
- Johnson-Laird, P. N. (1989). Mental models. In M. I. Posner (Ed.), *Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Kalish, C. W. (1995). Essentialism and graded membership in animal and artifact categories. *Memory & Cognition*, 23, 335-353.
- Keefe, D. E., McDaniel, M. A. (1993). The time course and durability of predictive inferences. *Journal of Memory and Language*, 32, 446-463.
- Keppel, G., & Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior*, 1, 153-161.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95, 163-182.
- Klatzky, R. L., Pellegrino, J. W., McCloskey, B. P., & Doherty, S. (1989). Can you squeeze a tomato? The role of motor representations in semantic sensibility judgments. *Journal of Memory and Language*, 28, 56-77.
- Kolers, P. A., & Roediger, H. L. III. (1984). Procedures of mind. *Journal of Verbal Learning and Verbal Behavior*, 23, 425-449.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago: University of Chicago Press.
- Lakoff, G. (1988). Smolensky, semantics, and the sensorimotor system. *Behavioral and Brain Sciences*, 11, 39-40.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lang, P. J. (1979). A bio-informational theory of emotional imagery. *Psychophysiology*, 16, 495-512.
- Langston, W., Kramer, D. C., & Glenberg, A. M. (1995). Mental models are not (very) spatial. Manuscript in preparation.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95, 492-527.
- Malter, A. J. (1996). An introduction to embodied cognition: Implications for consumer research. In K. P. Corfman & J. Lynch (Eds.), *Advances in consumer research* (Vol. 23). Provo, UT: Association for Consumer Research.
- Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 3-23.
- McClelland, J. L., Rumelhart, D. E., (1986). A distributed model of human learning and memory. In J. L. McClelland, D. E. Rumelhart, & the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2*.

- Psychological and biological models*. Cambridge, MA: MIT Press.
- McCutchen, D., & Perfetti, C. A. (1982). The visual tongue-twister: Phonological activation in silent reading. *Journal of Verbal Learning and Verbal Behavior*, 21, 627-687.
- McKoon, G., & Ratcliff, R. (1986a). Automatic activation of episodic information in a semantic memory task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 108-115.
- McKoon, G., & Ratcliff, R. (1986b). Inferences about predictable events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 82-91.
- McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, 99, 440-466.
- McKoon, G., Ratcliff, R., & Dell, G. S. (1986). A critical evaluation of the semantic-episodic distinction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 12, 1173-1190.
- McNamara, T. P. (1986). Mental representation of spatial relations. *Cognitive Psychology*, 18, 87-121.
- McNamara, T. P., Hardy, J. K., & Hirtle, S. C. (1989). Subjective hierarchies in spatial memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 211-227.
- Metcalfe, J. (1993). Novelty monitoring, metacognition, and control in a composite holographic associative recall model: implications for Korsakoff amnesia. *Psychological Review*, 100, 3-22.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, 90, 227-234.
- Montello, D. R., & Presson, C. C. (1993). Movement and orientation in surrounding and imaginal spaces. Manuscript in preparation.
- Morrow, D. G., Bower, G. H., & Greenspan, S. L. (1989). Updating situation models during narrative comprehension. *Journal of Verbal Learning and Verbal Behavior*, 28, 292-312.
- Murdock, B. B., Jr. (1974). *Human memory: Theory and data*. New York: Wiley.
- Murray, J. D., Klin, C. M., & Meyers, J. L. (1993). Forward inferences in narrative text. *Journal of Memory and Language*, 32, 464-473.
- Musen, G., & Squire, L. R. (1991). Normal acquisition of novel verbal information in amnesia. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17, 1095-1104.
- Muter, D. (1980). Very rapid forgetting. *Memory & Cognition*, 8, 174-179.
- Naveh-Benjamin, M., & Jonides, J. (1984). Maintenance rehearsal: A two-component analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 369-385.
- Needham, D. R., & Begg, I. M. (1991). Problem-oriented training promotes spontaneous analogical transfer: Memory-oriented training promotes memory for training. *Memory & Cognition*, 19, 543-557.
- Nelson, K. (1993). The psychological and social origins of autobiographical memory. *Psychological Science*, 4, 7-14.
- O'Brien, E. J., & Albrecht, J. E. (1992). Comprehension strategies in the development of a mental model. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18, 777-784.
- Oden, G. C. (1984). Integration of fuzzy linguistic information in language comprehension. *Fuzzy Sets and Systems*, 14, 29-41.
- Oden, G. C. (1987). Concept, knowledge, and thought. *Annual Review of Psychology*, 38, 203-227.
- Ohman, A., Fredrikson, M., Hugdahl, K., & Rimmo, P.-A. (1976). The premise of equipotentiality in human classical conditioning: Conditioned electrodermal responses to potentially phobic stimuli. *Journal of Experimental Psychology: General*, 105, 313-337.

- Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch and B. Lloyd (Eds.) *Cognition and categorization*. Hillsdale, NJ: Erlbaum.
- Plunkett, K., Sinha, C., Møller, M.F., & Strandsby, O. (1992). Symbol Grounding or the Emergence of Symbols? Vocabulary Growth in Children and a Connectionist Net. *Connection Science*, 4, 293-312.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353-363.
- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge: Cambridge University Press.
- Ratcliff, R., & McKoon, G. (1993, November). Bias in implicit memory tasks. Paper presented at the meeting of the Psychonomic Society, Washington, D.C.
- Ratcliff, R., McKoon, G., & Verwoerd, M. (1989). A bias interpretation of facilitation in perceptual identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 378-387.
- Regier, T. (1995). A model of the human capacity for categorizing spatial relations. *Cognitive Linguistics*, 6, 63-88.
- Reisberg, D., Rappaport, I., & O'Shaughnessy, M. (1984). Limits of working memory: The digit digit-span. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 10, 203-221.
- Rieser, J. J., Garing, A. E., & Young, M. F. (1994). Imagery, action and young children's spatial orientation: It's not being there that counts, it's what one has in mind. *Child Development*.
- Rinck, M., & Bower, G. H. (1995). Anaphora resolution and the focus of attention in situation models. *Journal of Memory and Language*, 34, 110-131.
- Roediger, H. L. III. (1990). Implicit memory: Retention without remembering. *American Psychologist*, 45, 1043-1056.
- Roediger, H. L. III., Guynn, M. J., & Jones, T. C. (1994). Implicit memory: A tutorial review. In G. d'Ydewelle, P. Eelen, and P. Bertelson (Eds.) *International perspectives on psychological science: The state of the art*. Vol 2. Hove, England: Lawrence Erlbaum.
- Rosch, E. (1973). Natural Categories. *Cognitive Psychology*, 4, 328-350.
- Rumelhart, D. E., Smolensky, P., McClelland, J.L., & Hinton, G.E (1986). Schemata and sequential thought processes in PDP models. In J.L. McClelland, D. E. Rumelhart, & the PDP Research Group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. Psychological and biological models*. Cambridge, MA: MIT Press.
- Saltz, E., & Donnenwerth-Nolan, S. (1981). Does motoric imagery facilitate memory for sentences? A selective interference test. *Journal of Verbal Learning and Verbal Behavior*, 20, 322-332.
- Sanford, A. J., & Moxey, L. M. (1995). Aspects of coherence in written language: A psychological perspective. In M. A. Gernsbacher & T. Givón, (Eds.), *Coherence in spontaneous text*. Philadelphia, PA: John Benjamins.
- Schacter, D. L., Cooper, L. A., & Delaney, S. M. (1990). Implicit memory for unfamiliar objects depends on access to structural descriptions. *Journal of Experimental Psychology: Human Learning and Memory*, 16, 5-24.
- Schneirla, T. C. (1959). An evolutionary and developmental theory of biphasic processes underlying approach and withdrawal. *Nebraska Symposium on Motivation*, 7, 1-42.
- Searle, J. R. (1980) Minds, brains and programs. *Behavioral and Brain Sciences*, 3, 417-424.
- Sharkey, A. J. C., & Sharkey, N. E. (1992). Weak contextual constraints in text and word priming. *Journal of Memory and Language*, 31, 543-572.
- Shelton, J. R., & Martin, R. C. (1992). How semantic is automatic semantic priming? *Journal of Experimental Psychology: Learning, Memory and Cognition*, 18, 1191-1210.
- Shepard, R. N. (1988). How fully should connectionism be activated? Two sources of excitation and one of inhibition. *Behavioral and Brain Sciences*, 11, 52.

- Shiffrar, M., Heptulla, S., O'Shaughnessy, M., and Freyd, J., (1993, November). *What does it mean to be sensitive to biological motion?* Paper presented at the meeting of the Psychonomic Society, Washington, D.C.
- Shulman, H. G. (1972). Semantic confusion errors in short-term memory. *Journal of Verbal Learning and Verbal Behavior*, 11, 221-227.
- Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 4, 592-604.
- Sloman, S. A., Hayman, C. A. G., Ohta, N., Law, J., & Tulving, E. (1988). Forgetting in primed fragment completion. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 14, 223-238.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12, 49-100.
- Tannenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., and Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 1632-1634.
- Taylor, H. A., & Tversky, B. (1992). Descriptions and depictions of environments. *Memory & Cognition*, 20, 483-496.
- Thorndike, E. L. (1932). *Fundamentals of learning*. New York: Teacher's College Press.
- Tulving, E. (1983). *Elements of episodic memory*. New York: Oxford University Press.
- Tulving, E. (1984). *Precis of Elements of episodic memory*. *The Behavioral and Brain Sciences*, 7, 223-268.
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, 247, 301-306.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, 113, 169-197.
- Van den Bergh, O., Vrana, S., & Eelen, P. (1990). Letters from the heart: Affective categorization of letter combinations in typists and nontypists. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 16, 1153-1161.
- Watkins, O. C., & Watkins, M. J. (1975). Build-up of proactive inhibition as a cue-overload effect. *Journal of Experimental Psychology: Human Learning and Memory*, 1, 442-452.
- Weldon, M. S., & Roediger, H. L. (1987). Altering retrieval demands reverses the picture superiority effect. *Memory & Cognition*, 15, 269-280.
- Wilson, S.G., Rinck, M., McNamara, T., Bower, G. H. & Morrow, D. G. (1993). Mental models and narrative comprehension: Some qualifications. *Journal of Memory and Learning*, 32, 141-154.
- Wilcoxon, H. C., Dragoin, W. B., & Kral, P. A. (1971). Illness-induced aversions in rat and quail: relative salience of visual and gustatory cues. *Science*, 171, 826-828.

Author Note

An important impetus for this paper was a pleasant conversation I had with Larry Barsalou while sitting on the steps of the Chicago Art Institute. Most of the ideas were developed further while I was a member of the research group "Mental Models in Discourse Processing" at the Zentrum für interdisziplinäre Forschung, Universität Bielefeld. I thank members of the group, and especially Tony Sanford, for many lively discussions. Larry Barsalou, William Epstein, Alan Garnham, and Morton Ann Gernsbacher in particular, provided much help in eliminating the most loopy ideas from an earlier draft. Additional valuable help was received from Jesse Prinz, Peter Gärdenfors, Henry Roediger, III, Robert Lockhart, G. Dorffner, and several anonymous reviewers. Preparation of this paper was supported in part by AFOSR AASERT Grant F49620-92-J-0310 and University of Wisconsin Graduate Research Committee Grant 940109 to Arthur Glenberg. Requests for reprints should be sent to Arthur Glenberg, Department of Psychology, University of Wisconsin, 1202 West Johnson Street, Madison, WI 53706, or by electronic mail to glenberg@facstaff.wisc.edu.